



ECS4150 Series

Software Release
v1.3.0.8

Web Management Guide

Web Management Guide

ECS4150-28T

L2+/L3 Lite Gigabit Ethernet Switch
with 24 1000BASE-T RJ-45 ports
and 4 10G SFP+ ports

ECS4150-28P

L2+/L3 Lite Gigabit Ethernet PoE Switch
with 24 1000BASE-T RJ-45 PoE ports
and 4 10G SFP+ ports

ECS4150-28F-I

L2+/L3 Lite Gigabit Ethernet Switch
with 24 100/1000M SFP ports
4 10/100/1000M RJ-45/SFP combo ports
and 4 1/10G SFP+ ports

ECS4150-54T

L2+/L3 Lite Gigabit Ethernet Switch
with 48 1000BASE-T RJ-45 ports
and 6 25G SFP28 ports

ECS4150-54P

L2+/L3 Lite Gigabit Ethernet PoE Switch
with 48 1000BASE-T RJ-45 PoE ports
and 6 25G SFP28 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:

Quick Start 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.

Documentation Notice This documentation is provided for general information purposes only. If any product feature details in this documentation conflict with the product datasheet, refer to the datasheet for the latest information.

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

Revision	Date	Change Description
v1.3.0.8	11/2023	Added: ◆ Support for ECS4150-28F-I ◆ "Connectivity Fault Management" on page 518 ◆ "Multicast VLAN Registration for IPv6" on page 645 ◆ "Excluded VLAN" on page 156 Updated: ◆ "Network Access (MAC Address Authentication)" on page 307
v1.2.0.6	06/2023	Added: ◆ Support for ECS4150-54T and ECS4150-54P Updated: ◆ "Port Configuration" on page 108 ◆ "Power over Ethernet" on page 438 ◆ "Configuring IGMP Snooping and Query Parameters" on page 578
v1.1.0.9	03/2023	Initial release

Contents

Configuring Remote Logon Authentication Servers 282

How to Use This Guide	3
Contents	5
Figures	17
Tables	33

Section I	Getting Started	35
	1 Introduction	37
	Key Features	37
	Description of Software Features	38
	System Defaults	43

Section II	Web Configuration	47
	2 Using the Web Interface	49
	Connecting to the Web Interface	49
	Navigating the Web Browser Interface	50
	Dashboard	50
	Configuration Options	51
	Panel Display	52
	Main Menu	53
	3 Basic Management Tasks	73
	Displaying System Information	74
	Displaying Hardware/Software Versions	75
	Configuring Support for Jumbo Frames	76
	Displaying Bridge Extension Capabilities	77

Managing System Files	78
Copying Files via FTP/FTPS/SFTP/SCP/TFTP or HTTP	78
Saving the Running Configuration to a Local File	81
Setting the Start-up File	82
Showing System Files	83
Automatic Operation Code Upgrade	83
Setting the System Clock	87
Setting the Time Manually	88
Setting the SNTP Polling Interval	89
Configuring NTP	89
Configuring Time Servers	90
Setting the Time Zone	94
Configuring Summer Time	95
Configuring the Console Port	97
Configuring Telnet Settings	99
Displaying CPU Utilization	100
Configuring CPU Guard	101
Displaying Memory Utilization	102
Resetting the System	103
4 Interface Configuration	107
Port Configuration	108
Configuring by Port List	108
Configuring by Port Range	110
Displaying Connection Status	110
Showing Port or Trunk Statistics	111
Displaying Statistical History	116
Displaying Transceiver Data	120
Configuring Transceiver Thresholds	121
Performing Cable Diagnostics	123
Trunk Configuration	125
Configuring a Static Trunk	126
Configuring a Dynamic Trunk	129
Displaying LACP Port Counters	135
Displaying LACP Settings and Status for the Local Side	136

Displaying LACP Settings and Status for the Remote Side	138
Configuring Load Balancing	139
Saving Power	141
Configuring Local Port Mirroring	142
Configuring Remote Port Mirroring	144
Sampling Traffic Flows	148
Configuring sFlow Receiver Settings	149
Configuring an sFlow Polling Instance	151
Traffic Segmentation	153
Enabling Traffic Segmentation	153
Configuring Uplink and Downlink Ports	154
Excluded VLAN	156
VLAN Trunking	158
5 VLAN Configuration	161
IEEE 802.1Q VLANs	161
Configuring VLAN Groups	165
Adding Static Members to VLANs	167
Configuring Dynamic VLAN Registration	171
IEEE 802.1Q Tunneling	175
Enabling QinQ Tunneling on the Switch	179
Creating CVLAN to SPVLAN Mapping Entries	180
Adding an Interface to a QinQ Tunnel	182
L2PT Tunneling	183
Configuring the L2PT Tunnel Address	185
Enabling L2PT for Selected Interfaces	186
Protocol VLANs	187
Configuring Protocol VLAN Groups	188
Mapping Protocol Groups to Interfaces	189
Configuring IP Subnet VLANs	191
Configuring MAC-based VLANs	193
Configuring VLAN Mirroring	195
Configuring VLAN Translation	197
6 Address Table Settings	201
Displaying the Dynamic Address Table	201

Clearing the Dynamic Address Table	202
Changing the Aging Time	203
Configuring MAC Address Learning	204
Setting Static Addresses	205
Configuring MAC Address Mirroring	207
Issuing MAC Address Traps	209
7 Spanning Tree Algorithm	211
Overview	211
Configuring Loopback Detection	213
Configuring Global Settings for STA	215
Displaying Global Settings for STA	221
Configuring Interface Settings for STA	222
Displaying Interface Settings for STA	226
Configuring Multiple Spanning Trees	229
Configuring Interface Settings for MSTP	233
8 Congestion Control	237
Rate Limiting	237
Storm Control	238
Automatic Traffic Control	240
Setting the ATC Timers	241
Configuring ATC Thresholds and Responses	243
9 Class of Service	247
Layer 2 Queue Settings	247
Setting the Default Priority for Interfaces	247
Selecting the Queue Mode	248
Mapping CoS Values to Egress Queues	251
Layer 3/4 Priority Settings	253
Setting Priority Processing to IP Precedence/DSCP or CoS	254
Mapping Ingress DSCP Values to Internal DSCP Values	255
Mapping CoS Priorities to Internal DSCP Values	257
Mapping IP Precedence Values to Internal DSCP Values	259
10 Quality of Service	263
Overview	263

Configuring a Class Map	264
Creating QoS Policies	267
Attaching a Policy Map to a Port	277
11 VoIP Traffic Configuration	279
Overview	279
Configuring VoIP Traffic	280
Configuring Telephony OUI	281
Configuring VoIP Traffic Ports	282
12 Security Measures	285
AAA (Authentication, Authorization and Accounting)	286
Configuring Local/Remote Logon Authentication	287
Configuring Remote Logon Authentication Servers	288
Configuring AAA Accounting	293
Configuring AAA Authorization	299
Configuring User Accounts	303
Web Authentication	305
Configuring Global Settings for Web Authentication	305
Configuring Interface Settings for Web Authentication	306
Network Access (MAC Address Authentication)	307
Configuring Global Settings for Network Access	310
Configuring Network Access for Ports	311
Configuring Port Link Detection	313
Configuring a MAC Address Filter	314
Displaying Secure MAC Address Information	316
Configuring HTTPS	317
Configuring Global Settings for HTTPS	317
Replacing the Default Secure-site Certificate	319
Configuring the Secure Shell	321
Configuring the SSH Server	323
Generating the Host Key Pair	324
Importing User Public Keys	325
Access Control Lists	327
Showing TCAM Utilization	328
Setting the ACL Name and Type	330

Configuring a Standard IPv4 ACL	332
Configuring an Extended IPv4 ACL	333
Configuring a Standard IPv6 ACL	335
Configuring an Extended IPv6 ACL	337
Configuring a MAC ACL	339
Configuring an ARP ACL	341
Binding a Port to an Access Control List	343
Showing ACL Hardware Counters	344
Filtering IP Addresses for Management Access	345
Configuring Port Security	348
Configuring 802.1X Port Authentication	350
Configuring 802.1X Global Settings	352
Configuring Port Authenticator Settings for 802.1X	353
Configuring Port Supplicant Settings for 802.1X	357
Displaying 802.1X Statistics	359
DoS Protection	362
DHCP Snooping	364
DHCP Snooping Global Configuration	367
DHCP Snooping VLAN Configuration	369
Configuring Ports for DHCP Snooping	370
Displaying DHCP Snooping Binding Information	371
DHCPv6 Snooping	372
DHCPv6 Snooping Global Configuration	374
DHCPv6 Snooping VLAN Configuration	376
Configuring Interfaces for DHCPv6 Snooping	377
Displaying DHCPv6 Snooping Binding Information	379
Displaying DHCPv6 Snooping Statistics	380
ND Snooping	381
ND Snooping Global Configuration	382
ND Snooping VLAN Configuration	384
Configuring Ports for ND Snooping	384
Displaying ND Snooping Binding Information	385
Displaying ND Snooping Prefix Information	386
IPv4 Source Guard	387
Configuring Ports for IPv4 Source Guard	387

Configuring Static Bindings for IPv4 Source Guard	389
Displaying Information for Dynamic IPv4 Source Guard Bindings	392
IPv6 Source Guard	393
Configuring Ports for IPv6 Source Guard	393
Configuring Static Bindings for IPv6 Source Guard	395
Displaying Information for Dynamic IPv6 Source Guard Bindings	398
ARP Inspection	399
Configuring Global Settings for ARP Inspection	400
Configuring VLAN Settings for ARP Inspection	402
Configuring Interface Settings for ARP Inspection	404
Displaying ARP Inspection Statistics	405
Displaying the ARP Inspection Log	406
Application Filter	407
13 Basic Administration Protocols	409
Configuring Event Logging	411
System Log Configuration	411
Remote Log Configuration	413
Sending Simple Mail Transfer Protocol Alerts	414
Link Layer Discovery Protocol	416
Setting LLDP Timing Attributes	416
Configuring LLDP Interface Attributes	418
Configuring LLDP Interface Civic-Address	422
Displaying LLDP Local Device Information	424
Displaying LLDP Remote Device Information	428
Displaying Device Statistics	436
Power over Ethernet	438
Setting the Switch's Overall PoE Power Budget	439
Setting the Port PoE Power Budget	440
Simple Network Management Protocol	442
Configuring Global Settings for SNMP	445
Setting Community Access Strings	445
Setting the Local Engine ID	447
Specifying a Remote Engine ID	448
Setting SNMPv3 Views	449

Configuring SNMPv3 Groups	452
Configuring Local SNMPv3 Users	456
Configuring Remote SNMPv3 Users	458
Specifying Trap Managers	461
Creating SNMP Notification Logs	465
Showing SNMP Statistics	467
Remote Monitoring	469
Configuring RMON Alarms	469
Configuring RMON Events	472
Configuring RMON History Samples	474
Configuring RMON Statistical Samples	477
Switch Clustering	479
Configuring General Settings for Clusters	480
Cluster Member Configuration	481
Managing Cluster Members	483
Setting a Time Range	484
Ethernet Ring Protection Switching	487
ERPS Global Configuration	491
ERPS VLAN Group Configuration	492
ERPS Ring Configuration	493
ERPS Instance Configuration	494
ERPS Forced and Manual Mode Operations	509
MLAG Configuration	513
MLAG Global Configuration	514
MLAG Domain Configuration	515
MLAG Group Configuration	516
Connectivity Fault Management	518
Configuring Global Settings for CFM	522
Configuring Interfaces for CFM	525
Configuring CFM Maintenance Domains	525
Configuring CFM Maintenance Associations	530
Configuring Maintenance End Points	534
Configuring Remote Maintenance End Points	536
Transmitting Link Trace Messages	538
Transmitting Loop Back Messages	540

Transmitting Delay-Measure Requests	542
Displaying Local MEPs	544
Displaying Details for Local MEPs	545
Displaying Local MIPs	547
Displaying Remote MEPs	548
Displaying Details for Remote MEPs	549
Displaying the Link Trace Cache	551
Displaying Fault Notification Settings	552
Displaying Continuity Check Errors	553
OAM Configuration	554
Enabling OAM on Local Ports	554
Displaying Statistics for OAM Messages	557
Displaying the OAM Event Log	557
Displaying the Status of Remote Interfaces	558
Configuring a Remote Loopback Test	559
Displaying Results of Remote Loopback Testing	561
UDLD Configuration	562
Configuring UDLD Protocol Intervals	563
Configuring UDLD Interface Settings	564
Displaying UDLD Neighbor Information	566
LBD Configuration	567
Configuring Global Settings for LBD	568
Configuring Interface Settings for LBD	570
Smart Pair Configuration	570
Configuring the Smart Pair Global Settings	571
Configuring Smart Pair Interface Settings	572
Show the Configured Smart Pair IDs	573
Display the Configured Smart Pair Port Members and Restore the Traffic	573
14 Multicast Filtering	575
Overview	575
Layer 2 IGMP (Snooping and Query for IPv4)	576
Configuring IGMP Snooping and Query Parameters	578
Specifying Static Interfaces for a Multicast Router	582
Assigning Interfaces to Multicast Services	584

Setting IGMP Snooping Status per Interface	586
Filtering IGMP Packets on an Interface	592
Displaying Multicast Groups Discovered by IGMP Snooping	593
Displaying IGMP Snooping Statistics	594
Filtering and Throttling IGMP Groups	599
Enabling IGMP Filtering and Throttling	600
Configuring IGMP Filter Profiles	600
Configuring IGMP Filtering and Throttling for Interfaces	603
MLD Snooping (Snooping and Query for IPv6)	604
Configuring MLD Snooping and Query Parameters	605
Setting Immediate Leave Status for MLD Snooping per Interface	607
Specifying Static Interfaces for an IPv6 Multicast Router	608
Assigning Interfaces to IPv6 Multicast Services	610
Filtering MLD Query Packets on an Interface	612
Showing MLD Snooping Groups and Source List	613
Displaying MLD Snooping Statistics	614
Filtering and Throttling MLD Groups	622
Enabling MLD Filtering and Throttling	623
Configuring MLD Filter Profiles	623
Configuring MLD Filtering and Throttling for Interfaces	626
Multicast VLAN Registration for IPv4	628
Configuring MVR Global Settings	629
Configuring MVR Domain Settings	631
Configuring MVR Group Address Profiles	633
Configuring MVR Interface Status	635
Assigning Static MVR Multicast Groups to Interfaces	638
Displaying MVR Receiver Groups	640
Displaying MVR Statistics	641
Multicast VLAN Registration for IPv6	645
Configuring MVR6 Global Settings	646
Configuring MVR6 Domain Settings	648
Configuring MVR6 Group Address Profiles	649
Configuring MVR6 Interface Status	652
Assigning Static MVR6 Multicast Groups to Interfaces	654
Displaying MVR6 Receiver Groups	656

Displaying MVR6 Statistics	657
15 IP Tools	663
Using the Ping Function	663
Using the Trace Route Function	665
Address Resolution Protocol	666
Basic ARP Configuration	667
Configuring Static ARP Addresses	668
Displaying Dynamic or Local ARP Entries	670
Displaying ARP Statistics	671
16 IP Configuration	673
Setting the Switch's IP Address (IP Version 4)	673
Configuring IPv4 Interface Settings	673
Setting the Switch's IP Address (IP Version 6)	677
Configuring the IPv6 Default Gateway	677
Configuring IPv6 Interface Settings	678
Configuring IPv6 Neighbor Addresses	683
Configuring an IPv6 Address	684
Configuring IPv6 ND Prefixes	687
Showing IPv6 Addresses	689
Showing the IPv6 Neighbor Cache	690
Showing IPv6 Statistics	692
Showing the MTU for Responding Destinations	697
17 General IP Routing	699
Overview	699
Initial Configuration	699
IP Routing and Switching	700
Routing Path Management	701
Routing Protocols	701
Configuring Static Routes	702
Displaying the Routing Table	703
18 IP Services	705
Domain Name Service	705
Configuring General DNS Service Parameters	705

Configuring a List of Domain Names	706
Configuring a List of Name Servers	708
Configuring Static DNS Host to Address Entries	709
Displaying the DNS Cache	710
Multicast Domain Name Service	711
Dynamic Host Configuration Protocol	712
Specifying a DHCP Client Identifier	713
Configuring DHCP L3 Relay Service	714
Configuring DHCP Relay Service with Option 82	716
Enabling DHCP Dynamic Provision	719
Configuring the DHCP Server	720
Configuring DHCPv6 Relay	728
Enabling DHCPv6 Dynamic Provision	730
Configuring the PPPoE Intermediate Agent	731
Configuring PPPoE IA Global Settings	731
Configuring PPPoE IA Interface Settings	733
Showing PPPoE IA Statistics	735

Section III	Appendices	737
A	Software Specifications	739
	Software Features	739
	Management Features	740
	Standards	741
	Management Information Bases	742
B	Troubleshooting	745
	Problems Accessing the Management Interface	745
	Using System Logs	746
C	License Information	747
	The GNU General Public License	747

Figures

Figure 1: Dashboard	50
Figure 2: System Information	74
Figure 3: General Switch Information	76
Figure 4: Configuring Support for Jumbo Frames	77
Figure 5: Displaying Bridge Extension Configuration	78
Figure 6: Copy Firmware	80
Figure 7: Saving the Running Configuration	82
Figure 8: Setting Start-Up Files	82
Figure 9: Displaying System Files	83
Figure 10: Configuring Automatic Code Upgrade	87
Figure 11: Manually Setting the System Clock	88
Figure 12: Setting the Polling Interval for SNTP	89
Figure 13: Configuring NTP	90
Figure 14: Specifying SNTP Time Servers	91
Figure 15: Adding an NTP Time Server	92
Figure 16: Showing the NTP Time Server List	92
Figure 17: Adding an NTP Authentication Key	93
Figure 18: Showing the NTP Authentication Key List	94
Figure 19: Setting the Time Zone	95
Figure 20: Configuring Summer Time	97
Figure 21: Console Port Settings	98
Figure 22: Telnet Connection Settings	100
Figure 23: Displaying CPU Utilization	101
Figure 24: Configuring CPU Guard	102
Figure 25: Displaying Memory Utilization	103
Figure 26: Restarting the Switch (Immediately)	105
Figure 27: Restarting the Switch (In)	105
Figure 28: Restarting the Switch (At)	106
Figure 29: Restarting the Switch (Regularly)	106

Figure 30: Configuring Connections by Port List	109
Figure 31: Configuring Connections by Port Range	110
Figure 32: Displaying Port Information	111
Figure 33: Showing Port Statistics (Table)	115
Figure 34: Showing Port Statistics (Chart)	116
Figure 35: Configuring a History Sample	118
Figure 36: Showing Entries for History Sampling	118
Figure 37: Showing Status of Statistical History Sample	119
Figure 38: Showing Current Statistics for a History Sample	119
Figure 39: Showing Ingress Statistics for a History Sample	120
Figure 40: Displaying Transceiver Data	121
Figure 41: Configuring Transceiver Thresholds	123
Figure 42: Performing Cable Tests	125
Figure 43: Configuring Static Trunks	126
Figure 44: Creating Static Trunks	127
Figure 45: Adding Static Trunks Members	128
Figure 46: Configuring Connection Parameters for a Static Trunk	128
Figure 47: Showing Information for Static Trunks	129
Figure 48: Configuring Dynamic Trunks	129
Figure 49: Configuring the LACP Aggregator Admin Key	132
Figure 50: Enabling LACP on a Port	133
Figure 51: Configuring LACP Parameters on a Port	133
Figure 52: Showing Members of a Dynamic Trunk	134
Figure 53: Configuring Connection Settings for a Dynamic Trunk	134
Figure 54: Showing Connection Parameters for Dynamic Trunks	135
Figure 55: Displaying LACP Port Counters	136
Figure 56: Displaying LACP Port Internal Information	137
Figure 57: Displaying LACP Port Remote Information	139
Figure 58: Configuring Load Balancing	140
Figure 59: Enabling Power Savings	142
Figure 60: Configuring Local Port Mirroring	142
Figure 61: Configuring Local Port Mirroring	143
Figure 62: Displaying Local Port Mirror Sessions	144
Figure 63: Configuring Remote Port Mirroring	144
Figure 64: Configuring Remote Port Mirroring (Source)	147

Figure 65: Configuring Remote Port Mirroring (Intermediate)	148
Figure 66: Configuring Remote Port Mirroring (Destination)	148
Figure 67: Configuring an sFlow Receiver	150
Figure 68: Showing sFlow Receivers	150
Figure 69: Configuring an sFlow Instance	152
Figure 70: Showing sFlow Instances	152
Figure 71: Enabling Traffic Segmentation	154
Figure 72: Configuring Members for Traffic Segmentation	155
Figure 73: Showing Traffic Segmentation Members	156
Figure 74: Configuring Excluded VLANs	158
Figure 75: Displaying Excluded VLANs	158
Figure 76: Configuring VLAN Trunking	159
Figure 77: Configuring VLAN Trunking	160
Figure 78: VLAN Compliant and VLAN Non-compliant Devices	163
Figure 79: Using GVRP	164
Figure 80: Creating Static VLANs	166
Figure 81: Modifying Settings for Static VLANs	166
Figure 82: Showing Static VLANs	167
Figure 83: Configuring Static Members by VLAN Index	170
Figure 84: Configuring Static VLAN Members by Interface	170
Figure 85: Configuring Static VLAN Members by Interface Range	171
Figure 86: Configuring Global Status of GVRP	173
Figure 87: Configuring GVRP for an Interface	173
Figure 88: Showing Dynamic VLANs Registered on the Switch	174
Figure 89: Showing the Members of a Dynamic VLAN	174
Figure 90: QinQ Operational Concept	176
Figure 91: Enabling QinQ Tunneling	180
Figure 92: Configuring CVLAN to SPVLAN Mapping Entries	181
Figure 93: Showing CVLAN to SPVLAN Mapping Entries	181
Figure 94: Adding an Interface to a QinQ Tunnel	183
Figure 95: Configuring the L2PT Tunnel Address	186
Figure 96: Enabling L2PT on Required Interfaces	187
Figure 97: Configuring Protocol VLANs	189
Figure 98: Displaying Protocol VLANs	189
Figure 99: Assigning Interfaces to Protocol VLANs	191

Figure 100: Showing the Interface to Protocol Group Mapping	191
Figure 101: Configuring IP Subnet VLANs	193
Figure 102: Showing IP Subnet VLANs	193
Figure 103: Configuring MAC-Based VLANs	195
Figure 104: Showing MAC-Based VLANs	195
Figure 105: Configuring VLAN Mirroring	196
Figure 106: Showing the VLANs to Mirror	197
Figure 107: Configuring VLAN Translation	197
Figure 108: Configuring VLAN Translation	198
Figure 109: Showing the Entries for VLAN Translation	199
Figure 110: Displaying the Dynamic MAC Address Table	202
Figure 111: Clearing Entries in the Dynamic MAC Address Table	203
Figure 112: Setting the Address Aging Time	204
Figure 113: Configuring MAC Address Learning	205
Figure 114: Configuring Static MAC Addresses	206
Figure 115: Displaying Static MAC Addresses	207
Figure 116: Mirroring Packets Based on the Source MAC Address	208
Figure 117: Showing the Source MAC Addresses to Mirror	208
Figure 118: Issuing MAC Address Traps (Global Configuration)	209
Figure 119: Issuing MAC Address Traps (Interface Configuration)	210
Figure 120: STP Root Ports and Designated Ports	212
Figure 121: MSTP Region, Internal Spanning Tree, Multiple Spanning Tree	212
Figure 122: Spanning Tree – Common Internal, Common, Internal	213
Figure 123: Configuring Port Loopback Detection	215
Figure 124: Configuring Global Settings for STA (STP)	219
Figure 125: Configuring Global Settings for STA (RSTP)	220
Figure 126: Configuring Global Settings for STA (MSTP)	220
Figure 127: Displaying Global Settings for STA	222
Figure 128: Determining the Root Port	223
Figure 129: Configuring Interface Settings for STA	226
Figure 130: STA Port Roles	228
Figure 131: Displaying Interface Settings for STA	229
Figure 132: Creating an MST Instance	230
Figure 133: Displaying MST Instances	231
Figure 134: Modifying the Priority for an MST Instance	231

Figure 135: Displaying Global Settings for an MST Instance	232
Figure 136: Adding a VLAN to an MST Instance	232
Figure 137: Displaying Members of an MST Instance	233
Figure 138: Configuring MSTP Interface Settings	234
Figure 139: Displaying MSTP Interface Settings	235
Figure 140: Configuring Rate Limits	238
Figure 141: Configuring Storm Control	240
Figure 142: Storm Control by Limiting the Traffic Rate	240
Figure 143: Storm Control by Shutting Down a Port	241
Figure 144: Configuring ATC Timers	243
Figure 145: Configuring ATC Interface Attributes	245
Figure 146: Setting the Default Port Priority	248
Figure 147: Setting the Queue Mode (Strict)	250
Figure 148: Setting the Queue Mode (WRR)	250
Figure 149: Setting the Queue Mode (Strict and WRR)	250
Figure 150: Mapping CoS Values to Egress Queues	252
Figure 151: Showing CoS Values to Egress Queue Mapping	253
Figure 152: Setting the Trust Mode	255
Figure 153: Configuring DSCP to DSCP Internal Mapping	256
Figure 154: Showing DSCP to DSCP Internal Mapping	257
Figure 155: Configuring CoS to DSCP Internal Mapping	259
Figure 156: Showing CoS to DSCP Internal Mapping	259
Figure 157: Configuring IP Precedence to DSCP Internal Mapping	261
Figure 158: Showing the IP Precedence to DSCP Internal Map	261
Figure 159: Configuring a Class Map	265
Figure 160: Showing Class Maps	266
Figure 161: Adding Rules to a Class Map	266
Figure 162: Showing the Rules for a Class Map	267
Figure 163: Configuring a Policy Map	274
Figure 164: Showing Policy Maps	275
Figure 165: Adding Rules to a Policy Map	276
Figure 166: Showing the Rules for a Policy Map	276
Figure 167: Attaching a Policy Map to a Port	277
Figure 168: Configuring a Voice VLAN	281
Figure 169: Configuring an OUI Telephony List	282

Figure 170: Showing an OUI Telephony List	282
Figure 171: Configuring Port Settings for a Voice VLAN	284
Figure 172: Configuring the Authentication Sequence	288
Figure 173: Authentication Server Operation	289
Figure 174: Configuring Remote Authentication Server (RADIUS)	292
Figure 175: Configuring Remote Authentication Server (TACACS+)	292
Figure 176: Configuring AAA Server Groups	293
Figure 177: Showing AAA Server Groups	293
Figure 178: Configuring Global Settings for AAA Accounting	296
Figure 179: Configuring AAA Accounting Methods	297
Figure 180: Showing AAA Accounting Methods	297
Figure 181: Configuring AAA Accounting Service for 802.1X Service	298
Figure 182: Configuring AAA Accounting Service for Command Service	298
Figure 183: Configuring AAA Accounting Service for Exec Service	298
Figure 184: Displaying a Summary of Applied AAA Accounting Methods	299
Figure 185: Displaying Statistics for AAA Accounting Sessions	299
Figure 186: Configuring AAA Authorization Methods	301
Figure 187: Showing AAA Authorization Methods	301
Figure 188: Configuring AAA Authorization Methods for Exec Service	302
Figure 189: Displaying the Applied AAA Authorization Method	302
Figure 190: Configuring User Accounts	304
Figure 191: Showing User Accounts	305
Figure 192: Configuring Global Settings for Web Authentication	306
Figure 193: Configuring Interface Settings for Web Authentication	307
Figure 194: Configuring Global Settings for Network Access	311
Figure 195: Configuring Interface Settings for Network Access	313
Figure 196: Configuring Link Detection for Network Access	314
Figure 197: Configuring a MAC Address Filter for Network Access	315
Figure 198: Showing the MAC Address Filter Table for Network Access	315
Figure 199: Showing Addresses Authenticated for Network Access	317
Figure 200: Configuring HTTPS	319
Figure 201: Downloading the Secure-Site Certificate	320
Figure 202: Configuring the SSH Server	324
Figure 203: Generating the SSH Host Key Pair	325
Figure 204: Showing the SSH Host Key Pair	325

Figure 205: Copying the SSH User's Public Key	326
Figure 206: Showing the SSH User's Public Key	327
Figure 207: Showing TCAM Utilization	330
Figure 208: Creating an ACL	331
Figure 209: Showing a List of ACLs	331
Figure 210: Configuring a Standard IPv4 ACL	333
Figure 211: Configuring an Extended IPv4 ACL	335
Figure 212: Configuring a Standard IPv6 ACL	337
Figure 213: Configuring an Extended IPv6 ACL	339
Figure 214: Configuring a MAC ACL	341
Figure 215: Configuring a ARP ACL	343
Figure 216: Binding a Port to an ACL	344
Figure 217: Showing ACL Statistics	345
Figure 218: Creating an IP Address Filter for Management Access	347
Figure 219: Showing IP Addresses Authorized for Management Access	347
Figure 220: Configuring Port Security	350
Figure 221: Configuring Port Authentication	351
Figure 222: Configuring Global Settings for 802.1X Port Authentication	353
Figure 223: Configuring Interface Settings for 802.1X Port Authenticator	357
Figure 224: Configuring Interface Settings for 802.1X Port Supplicant	359
Figure 225: Showing Statistics for 802.1X Port Authenticator	361
Figure 226: Showing Statistics for 802.1X Port Supplicant	362
Figure 227: Protecting Against DoS Attacks	364
Figure 228: Configuring Global Settings for DHCP Snooping	368
Figure 229: Configuring DHCP Snooping on a VLAN	369
Figure 230: Configuring the Port Mode for DHCP Snooping	371
Figure 231: Displaying the Binding Table for DHCP Snooping	372
Figure 232: Configuring Global Settings for DHCPv6 Snooping	376
Figure 233: Configuring DHCPv6 Snooping on a VLAN	377
Figure 234: Showing VLANs Enabled for DHCPv6 Snooping	377
Figure 235: Configuring the Trust Sate for DHCPv6 Snooping	379
Figure 236: Displaying the Binding Table for DHCPv6 Snooping	380
Figure 237: Displaying Statistics for DHCPv6 Snooping	381
Figure 238: ND Snooping Global Configuration	383
Figure 239: ND Snooping VLAN Configuration	384

Figure 240: ND Snooping Interface Configuration	385
Figure 241: Displaying the Binding List for ND Snooping	386
Figure 242: Displaying the Prefix List for ND Snooping	387
Figure 243: Setting the Filter Type for IPv4 Source Guard	389
Figure 244: Configuring Static Bindings for IPv4 Source Guard	391
Figure 245: Displaying Static Bindings for IPv4 Source Guard	392
Figure 246: Showing the IPv4 Source Guard Binding Table	393
Figure 247: Setting the Filter Type for IPv6 Source Guard	395
Figure 248: Configuring Static Bindings for IPv6 Source Guard	397
Figure 249: Displaying Static Bindings for IPv6 Source Guard	397
Figure 250: Showing the IPv6 Source Guard Binding Table	399
Figure 251: Configuring Global Settings for ARP Inspection	402
Figure 252: Configuring VLAN Settings for ARP Inspection	403
Figure 253: Configuring Interface Settings for ARP Inspection	404
Figure 254: Displaying Statistics for ARP Inspection	406
Figure 255: Displaying the ARP Inspection Log	407
Figure 256: Configuring Discarding or Forwarding of CDP/PVST Packets	407
Figure 257: Configuring Settings for System Memory Logs	412
Figure 258: Showing Error Messages Logged to System Memory	413
Figure 259: Configuring Settings for Remote Logging of Error Messages	414
Figure 260: Configuring SMTP Alert Messages	415
Figure 261: Configuring LLDP Timing Attributes	418
Figure 262: Configuring LLDP Interface Attributes	422
Figure 263: Configuring the Civic Address for an LLDP Interface	423
Figure 264: Showing the Civic Address for an LLDP Interface	424
Figure 265: Displaying Local Device Information for LLDP (General)	427
Figure 266: Displaying Local Device Information for LLDP (Port)	427
Figure 267: Displaying Local Device Information for LLDP (Port Details)	427
Figure 268: Displaying Remote Device Information for LLDP (Port)	434
Figure 269: Displaying Remote Device Information for LLDP (Port Details)	435
Figure 270: Displaying Remote Device Information for LLDP (End Node)	436
Figure 271: Displaying LLDP Device Statistics (General)	438
Figure 272: Displaying LLDP Device Statistics (Port)	438
Figure 273: Setting the Switch's PoE Budget	440
Figure 274: Setting a Port's PoE Budget	442

Figure 275: Configuring Global Settings for SNMP	445
Figure 276: Setting Community Access Strings	446
Figure 277: Showing Community Access Strings	446
Figure 278: Configuring the Local Engine ID for SNMP	447
Figure 279: Configuring a Remote Engine ID for SNMP	448
Figure 280: Showing Remote Engine IDs for SNMP	449
Figure 281: Creating an SNMP View	450
Figure 282: Showing SNMP Views	450
Figure 283: Adding an OID Subtree to an SNMP View	451
Figure 284: Showing the OID Subtree Configured for SNMP Views	451
Figure 285: Creating an SNMP Group	455
Figure 286: Showing SNMP Groups	455
Figure 287: Configuring Local SNMPv3 Users	457
Figure 288: Showing Local SNMPv3 Users	457
Figure 289: Changing a Local SNMPv3 User Group	458
Figure 290: Configuring Remote SNMPv3 Users	460
Figure 291: Showing Remote SNMPv3 Users	460
Figure 292: Configuring Trap Managers (SNMPv1)	464
Figure 293: Configuring Trap Managers (SNMPv2c)	464
Figure 294: Configuring Trap Managers (SNMPv3)	464
Figure 295: Showing Trap Managers	465
Figure 296: Creating SNMP Notification Logs	466
Figure 297: Showing SNMP Notification Logs	467
Figure 298: Showing SNMP Statistics	468
Figure 299: Configuring an RMON Alarm	471
Figure 300: Showing Configured RMON Alarms	471
Figure 301: Configuring an RMON Event	473
Figure 302: Showing Configured RMON Events	474
Figure 303: Configuring an RMON History Sample	475
Figure 304: Showing Configured RMON History Samples	476
Figure 305: Showing Collected RMON History Samples	476
Figure 306: Configuring an RMON Statistical Sample	478
Figure 307: Showing Configured RMON Statistical Samples	478
Figure 308: Showing Collected RMON Statistical Samples	479
Figure 309: Configuring a Switch Cluster	481

Figure 310: Configuring a Cluster Members	482
Figure 311: Showing Cluster Members	482
Figure 312: Showing Cluster Candidates	483
Figure 313: Managing a Cluster Member	484
Figure 314: Setting the Name of a Time Range	485
Figure 315: Showing a List of Time Ranges	485
Figure 316: Add a Rule to a Time Range	486
Figure 317: Showing the Rules Configured for a Time Range	486
Figure 318: ERPS Ring Components	488
Figure 319: Ring Interconnection Architecture (Multi-ring/Ladder Network)	489
Figure 320: Setting ERPS Global Status	492
Figure 321: Configuring ERPS VLAN Groups	493
Figure 322: Configuring an ERPS Ring	494
Figure 323: Sub-ring with Virtual Channel	503
Figure 324: Sub-ring without Virtual Channel	504
Figure 325: Creating an ERPS Instance	507
Figure 326: Configuring ERPS Instance Details	508
Figure 327: Showing Configured ERPS Instances	509
Figure 328: Blocking an ERPS Ring Port	513
Figure 329: MLAG Domain Topology	513
Figure 330: Configuring MLAG Global Status	515
Figure 331: Configuring MLAG Domains	516
Figure 332: MLAG Peer Operation	517
Figure 333: Configuring MLAG Groups	518
Figure 334: Single CFM Maintenance Domain	519
Figure 335: Multiple CFM Maintenance Domains	520
Figure 336: Configuring Global Settings for CFM	524
Figure 337: Configuring Interfaces for CFM	525
Figure 338: Configuring Maintenance Domains	529
Figure 339: Showing Maintenance Domains	529
Figure 340: Configuring Detailed Settings for Maintenance Domains	530
Figure 341: Creating Maintenance Associations	533
Figure 342: Showing Maintenance Associations	533
Figure 343: Configuring Detailed Settings for Maintenance Associations	534
Figure 344: Configuring Maintenance End Points	535

Figure 345: Showing Maintenance End Points	536
Figure 346: Configuring Remote Maintenance End Points	537
Figure 347: Showing Remote Maintenance End Points	538
Figure 348: Transmitting Link Trace Messages	540
Figure 349: Transmitting Loopback Messages	541
Figure 350: Transmitting Delay-Measure Messages	543
Figure 351: Showing Information on Local MEPs	544
Figure 352: Showing Detailed Information on Local MEPs	546
Figure 353: Showing Information on Local MIPs	547
Figure 354: Showing Information on Remote MEPs	548
Figure 355: Showing Detailed Information on Remote MEPs	550
Figure 356: Showing the Link Trace Cache	552
Figure 357: Showing Settings for the Fault Notification Generator	553
Figure 358: Showing Continuity Check Errors	554
Figure 359: Enabling OAM for Local Ports	556
Figure 360: Displaying Statistics for OAM Messages	557
Figure 361: Displaying the OAM Event Log	558
Figure 362: Displaying Status of Remote Interfaces	559
Figure 363: Running a Remote Loop Back Test	561
Figure 364: Displaying the Results of Remote Loop Back Testing	562
Figure 365: Configuring UDLD Protocol Intervals	564
Figure 366: Configuring UDLD Interface Settings	566
Figure 367: Displaying UDLD Neighbor Information	567
Figure 368: Configuring Global Settings for LBD	569
Figure 369: Configuring Interface Settings for LBD	570
Figure 370: Configuring the Smart Pair Global Settings (Adding a Smart Pair)	572
Figure 371: Configuring Interfaces for a Smart Pair	573
Figure 372: Displaying the Smart Pair IDs.	573
Figure 373: Displaying Port Members of a Smart Pair	574
Figure 374: Multicast Filtering Concept	576
Figure 375: Configuring General Settings for IGMP Snooping	582
Figure 376: Configuring a Static Interface for a Multicast Router	583
Figure 377: Showing Static Interfaces Attached a Multicast Router	584
Figure 378: Showing Current Interfaces Attached a Multicast Router	584
Figure 379: Assigning an Interface to a Multicast Service	585

Figure 380: Showing Static Interfaces Assigned to a Multicast Service	586
Figure 381: Configuring IGMP Snooping on a VLAN	591
Figure 382: Showing Interface Settings for IGMP Snooping	592
Figure 383: Dropping IGMP Query or Multicast Data Packets	593
Figure 384: Showing Multicast Groups Learned by IGMP Snooping	594
Figure 385: Displaying IGMP Snooping Statistics – Query	597
Figure 386: Displaying IGMP Snooping Statistics – VLAN	598
Figure 387: Displaying IGMP Snooping Statistics – Port	598
Figure 388: Displaying IGMP Snooping Statistics – Trunk	599
Figure 389: Enabling IGMP Filtering and Throttling	600
Figure 390: Creating an IGMP Filtering Profile	601
Figure 391: Showing the IGMP Filtering Profiles Created	602
Figure 392: Adding Multicast Groups to an IGMP Filtering Profile	602
Figure 393: Showing the Groups Assigned to an IGMP Filtering Profile	603
Figure 394: Configuring IGMP Filtering and Throttling Interface Settings	604
Figure 395: Configuring General Settings for MLD Snooping	607
Figure 396: Configuring Immediate Leave for MLD Snooping	608
Figure 397: Configuring a Static Interface for an IPv6 Multicast Router	609
Figure 398: Showing Static Interfaces Attached an IPv6 Multicast Router	609
Figure 399: Showing Current Interfaces Attached an IPv6 Multicast Router	610
Figure 400: Assigning an Interface to an IPv6 Multicast Service	611
Figure 401: Showing Static Interfaces Assigned to an IPv6 Multicast Service	611
Figure 402: Showing Current Interfaces Assigned to an IPv6 Multicast Service	612
Figure 403: Dropping MLD Query Packets	613
Figure 404: Showing IPv6 Multicast Services and Corresponding Sources	614
Figure 405: Displaying MLD Snooping Statistics – Input	618
Figure 406: Displaying MLD Snooping Statistics – Output	619
Figure 407: Displaying MLD Snooping Statistics – Query	619
Figure 408: Displaying MLD Snooping Statistics – Summary (Port/Trunk)	620
Figure 409: Displaying MLD Snooping Statistics – Summary (VLAN)	621
Figure 410: Clearing MLD Snooping Statistics	622
Figure 411: Enabling MLD Filtering and Throttling	623
Figure 412: Creating an MLD Filtering Profile	624
Figure 413: Showing the MLD Filtering Profiles Created	625
Figure 414: Adding Multicast Groups to an MLD Filtering Profile	625

Figure 415: Showing the Groups Assigned to an MLD Filtering Profile	626
Figure 416: Configuring MLD Filtering and Throttling Interface Settings	627
Figure 417: MVR Concept	628
Figure 418: Configuring Global Settings for MVR	631
Figure 419: Configuring Domain Settings for MVR	632
Figure 420: Configuring an MVR Group Address Profile	634
Figure 421: Displaying MVR Group Address Profiles	634
Figure 422: Assigning an MVR Group Address Profile to a Domain	635
Figure 423: Showing the MVR Group Address Profiles Assigned to a Domain	635
Figure 424: Configuring Interface Settings for MVR	638
Figure 425: Assigning Static MVR Groups to an Interface	639
Figure 426: Showing the Static MVR Groups Assigned to a Port	640
Figure 427: Displaying MVR Receiver Groups	641
Figure 428: Displaying MVR Statistics – Query	643
Figure 429: Displaying MVR Statistics – VLAN	644
Figure 430: Displaying MVR Statistics – Port	645
Figure 431: Configuring Global Settings for MVR6	648
Figure 432: Configuring Domain Settings for MVR6	649
Figure 433: Configuring an MVR6 Group Address Profile	650
Figure 434: Displaying MVR6 Group Address Profiles	651
Figure 435: Assigning an MVR6 Group Address Profile to a Domain	651
Figure 436: Showing MVR6 Group Address Profiles Assigned to a Domain	652
Figure 437: Configuring Interface Settings for MVR6	654
Figure 438: Assigning Static MVR6 Groups to a Port	655
Figure 439: Showing the Static MVR6 Groups Assigned to a Port	656
Figure 440: Displaying MVR6 Receiver Groups	657
Figure 441: Displaying MVR6 Statistics – Query	659
Figure 442: Displaying MVR6 Statistics – VLAN	660
Figure 443: Displaying MVR6 Statistics – Port	661
Figure 444: Displaying MVR6 Statistics – Trunk	662
Figure 445: Pinging a Network Device	664
Figure 446: Tracing the Route to a Network Device	666
Figure 447: Proxy ARP	667
Figure 448: Configuring General Settings for ARP	668
Figure 449: Configuring Static ARP Entries	669

Figure 450: Displaying Static ARP Entries	670
Figure 451: Displaying ARP Entries	670
Figure 452: Displaying ARP Statistics	671
Figure 453: Configuring a Static IPv4 Address	675
Figure 454: Configuring a Dynamic IPv4 Address	676
Figure 455: Showing the Configured IPv4 Address for an Interface	676
Figure 456: Configuring the IPv6 Default Gateway	678
Figure 457: Configuring General Settings for an IPv6 Interface	682
Figure 458: Configuring RA Guard for an IPv6 Interface	683
Figure 459: Configuring an IPv6 Neighbor Address	684
Figure 460: Configuring an IPv6 Address	687
Figure 461: Configuring IPv6 ND Prefixes	689
Figure 462: Showing Configured IPv6 Addresses	690
Figure 463: Showing IPv6 Neighbors	691
Figure 464: Showing IPv6 Statistics (IPv6)	696
Figure 465: Showing IPv6 Statistics (ICMPv6)	696
Figure 466: Showing IPv6 Statistics (UDP)	697
Figure 467: Showing Reported MTU Values	697
Figure 468: Virtual Interfaces and Layer 3 Routing	700
Figure 469: Configuring Static Routes	703
Figure 470: Displaying Static Routes	703
Figure 471: Displaying the Routing Table	704
Figure 472: Configuring General Settings for DNS	706
Figure 473: Configuring a List of Domain Names for DNS	707
Figure 474: Showing the List of Domain Names for DNS	707
Figure 475: Configuring a List of Name Servers for DNS	708
Figure 476: Showing the List of Name Servers for DNS	709
Figure 477: Configuring Static Entries in the DNS Table	710
Figure 478: Showing Static Entries in the DNS Table	710
Figure 479: Showing Entries in the DNS Cache	711
Figure 480: Configuring Multicast DNS	712
Figure 481: Specifying a DHCP Client Identifier	714
Figure 482: Layer 3 DHCP Relay Service	715
Figure 483: Configuring L3 DHCP Relay Service	716
Figure 484: Configuring DHCP Relay Information Option 82 Service	719

Figure 485: Enabling Dynamic Provisioning via DHCP	720
Figure 486: DHCP Server	721
Figure 487: Enabling the DHCP Server	722
Figure 488: Configuring Excluded Addresses on the DHCP Server	723
Figure 489: Showing Excluded Addresses on the DHCP Server	723
Figure 490: Configuring DHCP Server Address Pools (Network)	726
Figure 491: Configuring DHCP Server Address Pools (Host)	726
Figure 492: Showing Configured DHCP Server Address Pools	727
Figure 493: Shows Addresses Assigned by the DHCP Server	727
Figure 494: Enabling DHCPv6 Relay Agent for Unicast mode.	729
Figure 495: Enabling DHCPv6 Relay Agent for Multicast mode.	730
Figure 496: Enabling DHCPv6 Relay Agent for Multicast mode.	730
Figure 497: Enabling Dynamic Provisioning via DHCPv6	731
Figure 498: Configuring Global Settings for PPPoE Intermediate Agent	733
Figure 499: Configuring Interface Settings for PPPoE Intermediate Agent	735
Figure 500: Showing PPPoE Intermediate Agent Statistics	736

Tables

Table 1: Key Features	37
Table 2: System Defaults	43
Table 3: Web Page Configuration Buttons	51
Table 4: Switch Main Menu	53
Table 5: Predefined Summer-Time Parameters	96
Table 6: Port Statistics	112
Table 7: LACP Port Counters	135
Table 8: LACP Internal Configuration Information	136
Table 9: LACP Remote Device Configuration Information	138
Table 10: Traffic Segmentation Forwarding	154
Table 11: Recommended STA Path Cost Range	223
Table 12: Default STA Path Costs	223
Table 13: IEEE 802.1p Egress Queue Priority Mapping	251
Table 14: CoS Priority Levels	251
Table 15: Mapping Internal Per-hop Behavior to Hardware Queues	251
Table 16: Default Mapping of DSCP Values to Internal PHB/Drop Values	256
Table 17: Default Mapping of CoS/CFI to Internal PHB/Drop Precedence	258
Table 18: Mapping IP Precedence	260
Table 19: Default Mapping of IP Precedence to Internal PHB/Drop Values	260
Table 20: Dynamic QoS Profiles	309
Table 21: HTTPS System Support	318
Table 22: 802.1X Statistics	359
Table 23: ARP Inspection Statistics	405
Table 24: ARP Inspection Log	406
Table 25: Logging Levels	411
Table 26: LLDP MED Location CA Types	422
Table 27: Chassis ID Subtype	424
Table 28: System Capabilities	425
Table 29: Port ID Subtype	426

Table 30: Remote Port Auto-Negotiation Advertised Capability	429
Table 31: Maximum Number of Ports Providing Simultaneous Power	440
Table 32: SNMPv3 Security Models and Levels	443
Table 33: Supported Notification Messages	452
Table 34: ERPS Request/State Priority	510
Table 35: Remote MEP Priority Levels	527
Table 36: MEP Defect Descriptions	527
Table 37: OAM Operation State	555
Table 38: Remote Loopback Status	560
Table 39: Address Resolution Protocol	666
Table 40: ARP Statistics	671
Table 41: Show IPv6 Neighbors - display description	690
Table 42: Show IPv6 Statistics - display description	692
Table 43: Show MTU - display description	697
Table 44: Options 60, 66 and 67 Statements	713
Table 45: Options 55 and 124 Statements	713
Table 46: Troubleshooting Chart	745

Section I

Getting Started

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

This section includes these chapters:

["Introduction" on page 37](#)

Introduction

This switch provides a broad range of features for Layer 2 switching. 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/SFTP/SCP/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 1024 rules (global share), 64 ACLs, and a maximum of 1024 rules for an ACL
DHCP/DHCPv6	Client, Relay Option 82
DNS	Client and Proxy service
Port Configuration	Speed, duplex mode, and flow control
Port Trunking	Supports up to 28 trunks – static or dynamic trunking (LACP)
Port Mirroring	6 sessions, one or more source ports to one analysis port
Congestion Control	Rate Limiting Throttling for broadcast, multicast, unknown unicast storms
Address Table	16K MAC addresses in the forwarding table (shared for L2, multicast, Router IPv4/IPv6 host entries), 1K static MAC addresses, 2K L2 multicast groups

Table 1: Key Features (Continued)

Feature	Description
IP Version 4 and 6	Supports IPv4 and IPv6 addressing, and management
IEEE 802.1D Bridge	Supports dynamic data switching and addresses learning
Store-and-Forward Switching	Supported to ensure wire-speed switching while eliminating bad frames
Spanning Tree Algorithm	Supports standard STP, Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Trees (MSTP)
Virtual LANs	Up to 4094 using IEEE 802.1Q, 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)
Multicast Filtering	Supports IGMP snooping and query and Multicast VLAN Registration
Remote Device Management	Supports Ethernet OAM functions for attached CPEs (IEEE 802.3ah)

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 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 queuing ensures the minimum delay for moving real-time multimedia data across the network. While multicast filtering 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/SFTP/SCP/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. 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 Relay DHCP Relay is 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 28 trunks.

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 4 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 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.

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, and MLD Snooping and Query for IPv6 to manage multicast group registration. It also supports Multicast VLAN Registration (MVR for IPv4) 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	0 (disabled)

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	Disabled (Defaults: RSTP standard)
	Edge Ports	Disabled
LLDP	Status	Enabled
ERPS	Status	Disabled
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
	Ingress Port Priority	0
	Queue Mode	WRR
Traffic Prioritization	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
	Class of Service	Enabled
IP Settings	Management VLAN	VLAN 1
	IP Address	192.168.2.10
	Subnet Mask	255.255.255.0
	Default Gateway	Not configured
	DHCP	Client: Disabled
	DNS	Proxy service: Disabled
	BOOTP	Disabled

Table 2: System Defaults (Continued)

Function	Parameter	Default
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	Enabled
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

Section II

Web Configuration

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

This section includes these chapters:

- ◆ ["Using the Web Interface" on page 49](#)
- ◆ ["Basic Management Tasks" on page 73](#)
- ◆ ["Interface Configuration" on page 107](#)
- ◆ ["VLAN Configuration" on page 161](#)
- ◆ ["Address Table Settings" on page 201](#)
- ◆ ["Spanning Tree Algorithm" on page 211](#)
- ◆ ["Congestion Control" on page 237](#)
- ◆ ["Class of Service" on page 247](#)
- ◆ ["Quality of Service" on page 263](#)
- ◆ ["VoIP Traffic Configuration" on page 279](#)
- ◆ ["Security Measures" on page 285](#)
- ◆ ["Basic Administration Protocols" on page 409](#)
- ◆ ["Multicast Filtering" on page 575](#)
- ◆ ["IP Tools" on page 663](#)
- ◆ ["IP Configuration" on page 673](#)
- ◆ ["General IP Routing" on page 699](#)

◆ "IP Services" on page 705

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.



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. The default IP address and subnet mask for the switch is 192.168.2.10 and 255.255.255.0, with no default gateway. If this is not compatible with the subnet connected to the switch, you can configure it with a valid IP address, subnet mask, and default gateway. To configure this device as the default gateway, use the IP > Routing > Static Routes (Add) page, set the destination address to the required interface, and the next hop to null address 0.0.0.0.
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 303.](#))
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 222](#).

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 303](#) for more details.

Dashboard When your web browser connects with the switch’s web agent, the Dashboard is displayed as shown below. The Dashboard displays the main menu on the left side of the screen and System Information, CPU Utilization, Temperature, and Top 5 Most Active Interfaces on the right side. The main menu links are used to navigate to other menus, and display configuration parameters and statistics.

Figure 1: Dashboard





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."
	Saves current settings.

Table 3: Web Page Configuration Buttons (Continued)

Button	Action
	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).

Front Panel Indicators



Note: This manual covers the ECS4150-28T, ECS4150-28P, ECS4150-28F-I, ECS4150-54T, and ECS4150-54P Gigabit Ethernet switches. Other than the difference in the number and type of ports and the support for PoE on the ECS4150-28P and ECS4150-54P, there are no significant differences. Therefore, the screen display examples in this manual might be based on any of the switch models in this series.

Note: You can open a connection to the vendor's web site by clicking on the Edgecore logo.

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
Dashboard	Displays system information, CPU utilization, temperature, and top 5 most active interfaces.	50
System		
General	Provides basic system description, including contact information	74
Switch	Shows the number of ports, hardware version, power status, and firmware version numbers	75
Capability	Enables support for jumbo frames; shows the bridge extension parameters	76, 77
File		78
Copy	Allows the transfer and copying files	78
Automatic Operation Code Upgrade	Automatically upgrades operation code if a newer version is found on the server	83
Set Startup	Sets the startup file	82
Show	Shows the files stored in flash memory; allows deletion of files	83
Time		87
Configure General		
Manual	Manually sets the current time	88
SNTP	Configures SNTP polling interval	89
NTP	Configures NTP authentication parameters	89
Configure Time Server	Configures a list of SNTP servers	90
Configure SNTP Server	Sets the IP address for SNTP time servers	90
Add NTP Server	Adds NTP time server and index of authentication key	91
Show NTP Server	Shows list of configured NTP time servers	91
Add NTP Authentication Key	Adds key index and corresponding MD5 key	93
Show NTP Authentication Key	Shows list of configured authentication keys	93
Configure Time Zone	Sets the local time zone for the system clock	94
Configure Summer Time	Configures summer time settings	95
Console	Sets console port connection parameters	97
Telnet	Sets Telnet connection parameters	99
CPU Utilization	Displays information on CPU utilization	100
CPU Guard	Sets the CPU utilization watermark and threshold	101
Memory Status	Shows memory utilization parameters	102

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Reset	Restarts the switch immediately, at a specified time, after a specified delay, or at a periodic interval	103
Interface		107
Port		108
General		108
Configure by Port List	Configures connection settings per port	108
Configure by Port Range	Configures connection settings for a range of ports	110
Show Information	Displays port connection status	110
Statistics	Shows Interface, Etherlike, and RMON port statistics	111
Chart	Shows Interface, Etherlike, and RMON port statistics	111
History	Shows statistical history for specified interfaces	116
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	120 121
Cable Test	Performs cable diagnostics for selected port to diagnose any cable faults (short, open etc.) and report the cable length	123
Trunk		125
Static		126
Configure Trunk		126
Add	Creates a trunk, along with the first port member	126
Show	Shows the configured trunk identifiers	126
Add Member	Specifies ports to group into static trunks	126
Show Member	Shows the port members for the selected trunk	126
Configure General		126
Configure	Configures trunk connection settings	126
Show Information	Displays trunk connection settings	126
Dynamic		129
Configure Aggregator	Configures administration key and timeout for specific LACP groups	129
Configure Aggregation Port		126
Configure		126
General	Allows ports to dynamically join trunks	129
Actor	Configures parameters for link aggregation group members on the local side	129
Partner	Configures parameters for link aggregation group members on the remote side	129

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show Information		135
Counters	Displays statistics for LACP protocol messages	135
Internal	Displays configuration settings and operational state for the local side of a link aggregation	136
Neighbors	Displays configuration settings and operational state for the remote side of a link aggregation	138
Configure Trunk		129
Configure	Configures connection settings	129
Show	Displays port connection status	129
Show Member	Shows the active members in a trunk	129
Statistics	Shows Interface, Etherlike, and RMON port statistics	111
Chart	Shows Interface, Etherlike, and RMON port statistics	111
Load Balance	Sets the load-distribution method among ports in aggregated links	139
History	Shows statistical history for specified interfaces	116
Green Ethernet	Adjusts the power provided to ports based on the length of the cable used to connect to other devices	141
Mirror		142
Add	Sets the source and target ports for mirroring	142
Show	Shows the configured mirror sessions	142
RSPAN	Mirrors traffic from remote switches for analysis at a destination port on the local switch	144
sFlow	Configures flow sampling for receiver ports and instances	148
Configure Receiver	Creates an sFlow receiver on the switch	149
Configure Details	Enable an sFlow polling data source that polls periodically based on a specified time interval, or an sFlow data source instance that takes samples periodically based on the number of packets processed	151
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
Excluded VLAN		156
Add	Configures excluded VLAN sessions	156
Show	Displays the configured excluded VLAN sessions	156
VLAN Trunking	Allows unknown VLAN groups to pass through the specified interface	158
VLAN	Virtual LAN	161
Static		
Add	Creates VLAN groups	165

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show	Displays configured VLAN groups	165
Modify	Configures group name and administrative status	165
Edit Member by VLAN	Specifies VLAN attributes per VLAN	167
Edit Member by Interface	Specifies VLAN attributes per interface	167
Edit Member by Interface Range	Specifies VLAN attributes per interface range	167
Dynamic		171
Configure General	Enables GVRP VLAN registration protocol globally	171
Configure Interface	Configures GVRP status and timers per interface	171
Show Dynamic VLAN		171
Show VLAN	Shows the VLANs this switch has joined through GVRP	171
Show VLAN Member	Shows the interfaces assigned to a VLAN through GVRP	171
Tunnel	IEEE 802.1Q (QinQ) Tunneling	179
Configure Global	Sets tunnel mode for the switch	179
Configure Service	Sets a CVLAN to SPVLAN mapping entry	180
Configure Interface	Sets the tunnel mode for any participating interface	182
L2PT	Layer 2 Protocol Tunneling	183
Configure Global	Configures the destination MAC address for L2PT	185
Configure Interface	Enables L2PT on selected interfaces	186
Protocol		187
Configure Protocol		188
Add	Creates a protocol group, specifying supported protocols	188
Show	Shows configured protocol groups	188
Configure Interface		189
Add	Maps a protocol group to a VLAN	189
Show	Shows the protocol groups mapped to each VLAN	189
IP Subnet		191
Add	Maps IP subnet traffic to a VLAN	191
Show	Shows IP subnet to VLAN mapping	191
MAC-Based		193
Add	Maps traffic with specified source MAC address to a VLAN	193
Show	Shows source MAC address to VLAN mapping	193
Mirror	Configures VLAN traffic mirroring	195
Translation		197
Add	Maps VLAN IDs between the customer and service provider	197

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show	Displays the configuration settings for VLAN translation	197
MAC Address		201
Dynamic		
Configure Aging	Sets timeout for dynamically learned entries	203
Show Dynamic MAC	Displays dynamic entries in the address table	201
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	202
Learning Status	Enables MAC address learning on selected interfaces	204
Static	Configures static MAC addresses	205
Mirror	Configures MAC address mirroring	207
MAC Notification		209
Configure Global	Issues a trap when a dynamic MAC address is added or removed	209
Configure Interface	Enables MAC authentication traps on the current interface	209
Spanning Tree		211
Loopback Detection	Configures Loopback Detection parameters	213
STA	Spanning Tree Algorithm	
Configure Global		
Configure	Configures global bridge settings for STP, RSTP and MSTP	215
Show Information	Displays STA values used for the bridge	221
Configure Interface		
Configure	Configures interface settings for STA	222
Show Information	Displays interface settings for STA	226
MSTP	Multiple Spanning Tree Algorithm	229
Configure Global		229
Add	Configures initial VLAN and priority for an MST instance	229
Modify	Configures the priority or an MST instance	229
Show	Configures global settings for an MST instance	229
Add Member	Adds VLAN members for an MST instance	229
Show Member	Adds or deletes VLAN members for an MST instance	229
Show Information	Displays MSTP values used for the bridge	
Configure Interface		233
Configure	Configures interface settings for an MST instance	233
Show Information	Displays interface settings for an MST instance	233

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Traffic		
Rate Limit	Sets the input and output rate limits for a port	237
Storm Control	Sets the broadcast storm threshold for each interface	238
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	240
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	241
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	243
Priority		
Default Priority	Sets the default priority for each port or trunk	247
Queue	Sets queue mode for the switch; sets the service weight for each queue that will use a weighted or hybrid mode	248
Trust Mode	Selects DSCP or CoS priority processing	246
DSCP to DSCP	Maps DSCP values in incoming packets to per-hop behavior and drop precedence values for internal priority processing	255
Cos to DSCP	Maps CoS/CFI values in incoming packets to per-hop behavior and drop precedence values for priority processing	257
IP Precedence to DSCP	Maps IP Precedence values to DSCP values	259
PHB to Queue	Maps internal per-hop behavior values to hardware queues	251
DiffServ		263
Configure Class		264
Add	Creates a class map for a type of traffic	264
Show	Shows configured class maps	264
Modify	Modifies the name of a class map	264
Add Rule	Configures the criteria used to classify ingress traffic	264
Show Rule	Shows the traffic classification rules for a class map	264
Configure Policy		255
Add	Creates a policy map to apply to multiple interfaces	255
Show	Shows configured policy maps	255
Modify	Modifies the name of a policy map	255
Add Rule	Sets the boundary parameters used for monitoring inbound traffic, and the action to take for conforming and non-conforming traffic	255
Show Rule	Shows the rules used to enforce bandwidth policing for a policy map	255
Configure Interface	Applies a policy map to an ingress port	259
VoIP	Voice over IP	279

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Configure Global	Configures auto-detection of VoIP traffic, sets the Voice VLAN, and VLAN aging time	280
Configure OUI		281
Add	Maps the OUI in the source MAC address of ingress packets to the VoIP device manufacturer	281
Show	Shows the OUI telephony list	281
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	282
Security		285
AAA	Authentication, Authorization and Accounting	286
System Authentication	Configures authentication sequence – local, RADIUS, and TACACS	287
Server		288
Configure Server	Configures RADIUS and TACACS server message exchange settings	288
Configure Group		288
Add	Specifies a group of authentication servers and sets the priority sequence	288
Show	Shows the authentication server groups and priority sequence	288
Accounting	Enables accounting of requested services for billing or security purposes	293
Configure Global	Specifies the interval at which the local accounting service updates information to the accounting server	293
Configure Method		293
Add	Configures accounting for various service types	293
Show	Shows the accounting settings used for various service types	293
Configure Service	Sets the accounting method applied to specific interfaces for 802.1X, CLI command privilege levels for the console port, and for Telnet	293
Show Information		293
Summary	Shows the configured accounting methods, and the methods applied to specific interfaces	293
Statistics	Shows basic accounting information recorded for user sessions	293
Authorization	Enables authorization of requested services	299
Configure Method		299
Add	Configures authorization for various service types	299
Show	Shows the authorization settings used for various service types	299
Configure Service	Sets the authorization method applied used for the console port, and for Telnet	299
Show Information	Shows the configured authorization methods, and the methods applied to specific interfaces	299

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
User Accounts		303
Add	Configures user names, passwords, and access levels	303
Show	Shows authorized users	303
Modify	Modifies user attributes	303
Web Authentication	Allows authentication and access to the network when 802.1X or Network Access authentication are infeasible or impractical	305
Configure Global	Configures general protocol settings	305
Configure Interface	Enables Web Authentication for individual ports	306
Network Access	MAC address-based network access authentication	307
Configure Global	Enables aging for authenticated MAC addresses, and sets the time period after which a connected MAC address must be reauthenticated	310
Configure Interface		311
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	311
Configure MAC Filter		314
Add	Specifies MAC addresses exempt from authentication	314
Show	Shows the list of exempt MAC addresses	314
Show Information	Shows the authenticated MAC address list	316
HTTPS	Secure HTTP	317
Configure Global	Enables HTTPS, and specifies the UDP port to use	317
Copy Certificate	Replaces the default secure-site certificate	319
SSH	Secure Shell	321
Configure Global	Configures SSH server settings	323
Configure Host Key		324
Generate	Generates the host key pair (public and private)	324
Show	Displays RSA host keys; deletes host keys	324
Configure User Key		325
Copy	Imports user public keys from a TFTP server	325
Show	Displays RSA user keys; deletes user keys	325
ACL	Access Control Lists	327
Configure ACL		330
Show TCAM	Shows utilization parameters for TCAM	328
Add	Adds an ACL based on IP or MAC address filtering	330
Show	Shows the name and type of configured ACLs	330
Add Rule	Configures packet filtering based on IP or MAC addresses and other packet attributes	330

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show Rule	Shows the rules specified for an ACL	330
Configure Interface	Binds a port to the specified ACL and time range	
Configure	Binds a port to the specified ACL and time range	343
Show Hardware Counters	Shows statistics for ACL hardware counters	344
IP Filter		345
Add	Sets IP addresses of clients allowed management access via the web, SNMP, and Telnet	345
Show	Shows the addresses to be allowed management access	345
Port Security	Configures per port security, including status, response for security breach, and maximum allowed MAC addresses	348
Port Authentication	IEEE 802.1X	350
Configure Global	Enables authentication and EAPOL pass-through	352
Configure Interface	Sets authentication parameters for individual ports	353
Show Statistics	Displays protocol statistics for the selected port	359
DoS Protection	Protects against Denial-of-Service attacks	362
DHCP Snooping		364
Configure Global	Enables DHCP snooping globally, MAC-address verification, information option; and sets the information policy	367
Configure VLAN	Enables DHCP snooping on a VLAN	369
Configure Interface	Sets the trust mode for an interface	370
Show Information	Displays the DHCP Snooping binding information	371
DHCP Snooping6		364
Configure Global	Enables DHCPv6 snooping globally, MAC-address verification, information option; and sets the information policy	367
Configure VLAN	Enables DHCPv6 snooping on a VLAN	369
Configure Interface	Sets the trust mode for an interface	370
Show Information	Displays the DHCPv6 Snooping binding information	371
ND Snooping		381
Configure Global	Enables ND Snooping globally on the switch	382
Configure VLAN	Enables ND Snooping on a VLAN or range of VLANs	384
Configure Interface	Sets the trust mode for an interface	384
Show Information	Displays the ND Snooping binding and prefix information	385
IP Source Guard	Filters IP traffic based on static entries in the IP Source Guard table, or dynamic entries in the DHCP Snooping table	387
General	Enables IP source guard and selects filter type per port	387
Static Binding		389
Configure ACL Table		389

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Add	Adds static addresses to the source guard ACL binding table	389
Show	Shows static addresses in the source guard ACL binding table	389
Configure MAC Table		389
Add	Adds static addresses to the source guard MAC address binding table	389
Show	Shows static addresses in the source guard MAC address binding table	389
Dynamic Binding	Displays the source-guard binding table for a selected interface	392
IPv6 Source Guard	Filters IPv6 traffic based on static entries in the IP Source Guard table, or dynamic entries in the DHCP Snooping table	393
Port Configuration	Enables IPv6 source guard and selects filter type per port	393
Static Binding		395
Add	Adds a static addresses to the source-guard binding table	395
Show	Shows static addresses in the source-guard binding table	395
Dynamic Binding	Displays the source-guard binding table for a selected interface	398
ARP Inspection		399
Configure General	Enables inspection globally, configures validation of additional address components, and sets the log rate for packet inspection	400
Configure VLAN	Enables ARP inspection on specified VLANs	402
Configure Interface	Sets the trust mode for ports, and sets the rate limit for packet inspection	404
Show Information		405
Show Statistics	Displays statistics on the inspection process	405
Show Log	Shows the inspection log list	406
Application Filter	Discards CDP or PVST packets	407
Administration		409
Log		411
System		411
Configure Global	Stores error messages in local memory	411
Show System Logs	Shows logged error messages	411
Remote	Configures the logging of messages to a remote logging process	413
SMTP	Sends an SMTP client message to a participating server	414
LLDP		416
Configure Global	Configures global LLDP timing parameters	416
Configure Interface		418
Configure General	Sets the message transmission mode; enables SNMP notification; and sets the LLDP attributes to advertise	418
Add CA-Type	Specifies the physical location of the device attached to an interface	422

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show Local Device Information		424
General	Displays general information about the local device	424
Port/Trunk	Displays information about each interface	424
Show Remote Device Information		428
Port/Trunk	Displays information about a remote device connected to a port on this switch	428
Port/Trunk Details	Displays detailed information about a remote device connected to this switch	428
Show Device Statistics		436
General	Displays statistics for all connected remote devices	436
Port/Trunk	Displays statistics for remote devices on a selected port or trunk	436
PoE*	Power over Ethernet	438
PSE	Displays the PoE power status for the switch	440
SNMP	Simple Network Management Protocol	442
Configure Global	Enables SNMP agent status, and sets related trap functions	445
Configure Engine		447
Set Engine ID	Sets the SNMP v3 engine ID on this switch	447
Add Remote Engine	Sets the SNMP v3 engine ID for a remote device	448
Show Remote Engine	Shows configured engine ID for remote devices	448
Configure View		449
Add View	Adds an SNMP v3 view of the OID MIB	449
Show View	Shows configured SNMP v3 views	449
Add OID Subtree	Specifies a part of the subtree for the selected view	449
Show OID Subtree	Shows the subtrees assigned to each view	449
Configure Group		452
Add	Adds a group with access policies for assigned users	452
Show	Shows configured groups and access policies	452
Configure User		
Add Community	Configures community strings and access mode	456
Show Community	Shows community strings and access mode	456
Add SNMPv3 Local User	Configures SNMPv3 users on this switch	456
Show SNMPv3 Local User	Shows SNMPv3 users configured on this switch	456
Change SNMPv3 Local User Group	Assign a local user to a new group	456
Add SNMPv3 Remote User	Configures SNMPv3 users from a remote device	458
Show SNMPv3 Remote User	Shows SNMPv3 users set from a remote device	456

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Configure Trap		461
Add	Configures trap managers to receive messages on key events that occur on this switch	461
Show	Shows configured trap managers	461
Configure Notify Filter		
Add	Creates an SNMP notification log	465
Show	Shows the configured notification logs	465
Show Statistics	Shows the status of SNMP communications	467
RMON	Remote Monitoring	469
Configure Global		
Add		
Alarm	Sets threshold bounds for a monitored variable	469
Event	Creates a response event for an alarm	472
Show		469
Alarm	Shows all configured alarms	469
Event	Shows all configured events	472
Configure Interface		
Add		
History	Periodically samples statistics on a physical interface	474
Statistics	Enables collection of statistics on a physical interface	477
Show		
History	Shows sampling parameters for each entry in the history group	474
Statistics	Shows sampling parameters for each entry in the statistics group	477
Show Details		
History	Shows sampled data for each entry in the history group	474
Statistics	Shows sampled data for each entry in the history group	477
Cluster		479
Configure Global	Globally enables clustering for the switch; sets Commander status	480
Configure Member	Adds switch Members to the cluster	481
Show Member	Shows cluster switch member; managed switch members	483
Time Range	Configures the time to apply an ACL	484
Add	Specifies the name of a time range	484
Show	Shows the name of configured time ranges	484
Add Rule		484
Absolute	Sets exact time or time range	484

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Periodic	Sets a recurrent time	484
Show Rule	Shows the time specified by a rule	484
ERPS	Ethernet Ring Protection Switching	487
Configure Global	Activates ERPS globally	491
Configure Domain		494
Add	Creates an ERPS ring	494
Show	Shows list of configured ERPS rings, status, and settings	494
Configure Details	Configures ring parameters	494
Configure Operation	Blocks a ring port using Forced Switch or Manual Switch commands	509
MLAG		513
Configure Global	Enables MLAG globally	514
Configure Domain	Configures MLAG domains and peer links	515
Configure Group	Configures MLAG member interfaces	516
CFM	Connectivity Fault Management	518
Configure Global	Configures global settings, including administrative status, cross-check start delay, link trace, and SNMP traps	522
Configure Interface	Configures administrative status on an interface	525
Configure MD	Configure Maintenance Domains	525
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	525
Configure Details	Configures the archive hold time and fault notification settings	525
Show	Shows list of configured maintenance domains	525
Configure MA	Configure Maintenance Associations	530
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	530
Configure Details	Configures detailed settings, including continuity check status and interval level, cross-check status, and alarm indication signal parameters	530
Show	Shows list of configured maintenance associations	530
Configure MEP	Configures Maintenance End Points	534
Add	Configures MEPs at the domain boundary to provide management access for each maintenance association	534
Show	Shows list of configured maintenance end points	534
Configure Remote MEP	Configures Remote Maintenance End Points	536
Add	Configures a static list of remote MEPs for comparison against the MEPs learned through continuity check messages	536

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show	Shows list of configured remote maintenance end points	536
Transmit Link Trace	Sends link trace messages to isolate connectivity faults by tracing the path through a network to the designated target node	538
Transmit Loopback	Sends loopback messages to isolate connectivity faults by requesting a target node to echo the message back to the source	540
Transmit Delay Measure	Sends periodic delay-measure requests to a specified MEP within a maintenance association	542
Show Information		
Show Local MEP	Shows the MEPs configured on this device	544
Show Local MEP Details	Displays detailed CFM information about a specified local MEP in the continuity check database	545
Show Local MIP	Shows the MIPs on this device discovered by the CFM protocol	547
Show Remote MEP	Shows MEPs located on other devices which have been discovered through continuity check messages, or statically configured in the MEP database	548
Show Remote MEP Details	Displays detailed CFM information about a specified remote MEP in the continuity check database	549
Show Link Trace Cache	Shows information about link trace operations launched from this device	551
Show Fault Notification Generator	Displays configuration settings for the fault notification generator	552
Show Continuity Check Error	Displays CFM continuity check errors logged on this device	553
OAM	Operation, Administration, and Maintenance	554
Interface	Enables OAM on specified port, sets the mode to active or passive, and enables the reporting of critical events or errored frame events	554
Counters	Displays statistics on OAM PDUs	557
Event Log	Displays the log for recorded link events	557
Remote Interface	Displays information about attached OAM-enabled devices	558
Remote Loopback	Performs a loopback test on the specified port	559
UDLD	UniDirectional Link Detection	562
Configure Global	Configures the message probe interval, detection interval, and recovery interval	563
Configure Interface	Enables UDLD and aggressive mode which reduces the shut-down delay after loss of bidirectional connectivity is detected	564
Show Information	Displays UDLD neighbor information, including neighbor state, expiration time, and protocol intervals	566
LBD	Loopback Detection	567
Configure Global	Enables loopback detection globally, specifies the interval at which to transmit control frames, specifies the interval to wait before releasing an interface from shutdown state, specifies response to detect loopback, and traps to send	568
Configure Interface	Enables loopback detection per interface	570

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Tools		
Ping	Sends ICMP echo request packets to another node on the network	463
Trace Route	Shows the route packets take to the specified destination	464
ARP	Shows entries in the Address Resolution Protocol cache	466
IP		673
General		
Routing Interface		
Add Address	Configures an IP interface for a VLAN	673
Show Address	Shows the IP interfaces assigned to a VLAN	673
Routing		
Static Routes		702
Add	Configures static routing entries	702
Show	Shows static routing entries	702
Routing Table	Shows all routing entries, including local, static and dynamic routes	703
IPv6 Configuration		677
Configure Global	Sets an IPv6 default gateway for traffic with no known next hop	677
Configure Interface	Configures IPv6 interface address using auto-configuration or link-local address, and sets related protocol settings	678
Configure Neighbor	Configures static IPv6 neighbor addresses	683
Add IPv6 Address	Adds an global unicast, EUI-64, or link-local IPv6 address to an interface	684
Configure IPv6 ND Prefix	Configures IPv6 prefixes to include in router advertisements	687
Show IPv6 Address	Show the IPv6 addresses assigned to an interface	689
Show IPv6 Neighbor Cache	Displays information in the IPv6 neighbor discovery cache	690
Show Statistics		692
IPv6	Shows statistics about IPv6 traffic	692
ICMPv6	Shows statistics about ICMPv6 messages	692
UDP	Shows statistics about UDP messages	692
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	697
IP Service		705
DNS	Domain Name Service	
General		705
Configure Global	Enables DNS lookup; defines the default domain name appended to incomplete host names	705

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	706
Show Domain Names	Shows the configured domain name list	706
Add Name Server	Specifies IP address of name servers for dynamic lookup	708
Show Name Servers	Shows the name server address list	708
Static Host Table	Manually configures static entries in the DNS table	709
Add	Configures static entries for domain name to address mapping	709
Show	Shows the list of static mapping entries	709
Modify	Modifies the static address mapped to the selected host name	709
Cache	Displays cache entries discovered by designated name servers	710
Multicast DNS	Configures multicast DNS lookup on the local network without the need for a dedicated server	705
DHCP	Dynamic Host Configuration Protocol	712
Client	Specifies the DHCP client identifier for an interface	713
L3 Relay	Specifies DHCP relay servers	714
Relay Option 82	Enables and configures DHCP relay with Option 82 information	716
Dynamic Provision	Enables dynamic provisioning via DHCP	719
Server	Configures a DHCP service on the switch	720
DHCPv6	Dynamic Host Configuration Protocol for IPv6	728
Relay	Specifies DHCPv6 relay servers	728
Dynamic Provision	Enables dynamic provisioning via DHCPv6	730
PPPoE Intermediate Agent	Point-to-Point Protocol over Ethernet Intermediate Agent	720
Multicast		575
IGMP Snooping		576
General	Enables multicast filtering; configures parameters for multicast snooping	578
Multicast Router		582
Add Static Multicast Router	Assigns ports that are attached to a neighboring multicast router	582
Show Static Multicast Router	Displays ports statically configured as attached to a neighboring multicast router	582
Show Current Multicast Router	Displays ports attached to a neighboring multicast router, either through static or dynamic configuration	582
IGMP Member		584
Add Static Member	Statically assigns multicast addresses to the selected VLAN	584
Show Static Member	Shows multicast addresses statically configured on the selected VLAN	584

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Interface		586
Configure VLAN	Configures IGMP snooping per VLAN interface	586
Show VLAN Information	Shows IGMP snooping settings per VLAN interface	586
Configure Port	Configures the interface to drop IGMP query packets or all multicast data packets	592
Configure Trunk	Configures the interface to drop IGMP query packets or all multicast data packets	592
Forwarding Entry	Displays the current multicast groups learned through IGMP Snooping	593
Filter		599
Configure General	Enables IGMP filtering for the switch	600
Configure Profile		600
Add	Adds IGMP filter profile; and sets access mode	600
Show	Shows configured IGMP filter profiles	600
Add Multicast Group Range	Assigns multicast groups to selected profile	600
Show Multicast Group Range	Shows multicast groups assigned to a profile	600
Configure Interface	Assigns IGMP filter profiles to port interfaces and sets throttling action	603
Statistics		594
Show Query Statistics	Shows statistics for query-related messages	594
Show VLAN Statistics	Shows statistics for protocol messages, number of active groups	594
Show Port Statistics	Shows statistics for protocol messages, number of active groups	594
Show Trunk Statistics	Shows statistics for protocol messages, number of active groups	594
MLD Snooping		604
General	Enables multicast filtering; configures parameters for IPv6 multicast snooping	605
Interface	Configures Immediate Leave status for a VLAN	607
Multicast Router		608
Add Static Multicast Router	Assigns ports that are attached to a neighboring multicast router	608
Show Static Multicast Router	Displays ports statically configured as attached to a neighboring multicast router	608
Show Current Multicast Router	Displays ports attached to a neighboring multicast router, either through static or dynamic configuration	608
MLD Member		610
Add Static Member	Statically assigns multicast addresses to the selected VLAN	610
Show Static Member	Shows multicast addresses statically configured on the selected VLAN	610
Show Current Member	Shows multicast addresses associated with the selected VLAN, either through static or dynamic configuration	610

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Filter		622
Configure General	Enables MLD filtering for the switch	623
Configure Profile		623
Add	Adds MLD filter profile; and sets access mode	623
Show	Shows configured MLD filter profiles	623
Add Multicast Group Range	Assigns multicast groups to selected profile	623
Show Multicast Group Range	Shows multicast groups assigned to a profile	623
Query Drop	Configures the interface to drop MLD query packets	612
Group Information	Displays known multicast groups, member ports, the means by which each group was learned, and the corresponding source list	613
Statistics		614
Input	Shows statistics for MLD ingress traffic	610
Output	Shows statistics for MLD egress traffic	610
Query	Shows statistics for query-related messages	610
Summary	Shows summary statistics for querier and report/leave messages	610
Clear	Clears all MLD statics or statistics for specified VLAN/port	610
MVR	Multicast VLAN Registration	628
Configure Global	Configures proxy switching and robustness value	629
Configure Domain	Enables MVR for a domain, sets the MVR VLAN, forwarding priority, and upstream source IP	631
Configure Profile		633
Add	Configures multicast stream addresses	633
Show	Shows multicast stream addresses	633
Associate Profile		633
Add	Maps an address profile to a domain	633
Show	Shows addresses profile to domain mapping	633
Configure Interface	Configures MVR interface type and immediate leave mode; also displays MVR operational and active status	635
Configure Static Group Member		638
Add	Statically assigns MVR multicast streams to an interface	638
Show	Shows MVR multicast streams assigned to an interface	638
Show Member	Shows the multicast groups assigned to an MVR VLAN, the source address of the multicast services, and the interfaces with active subscribers	640
Show Statistics		641
Show Query Statistics	Shows statistics for query-related messages	641
Show VLAN Statistics	Shows statistics for protocol messages and number of active groups	641

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show Port Statistics	Shows statistics for protocol messages and number of active groups	641
Show Trunk Statistics	Shows statistics for protocol messages and number of active groups	641
MVR6	Multicast VLAN Registration for IPv6	645
Configure Global	Configures proxy switching and robustness value	646
Configure Domain	Enables MVR for a domain, sets the MVR VLAN, forwarding priority, and upstream source IP	648
Configure Profile		649
Add	Configures multicast stream addresses	649
Show	Shows multicast stream addresses	649
Associate Profile		649
Add	Maps an address profile to a domain	649
Show	Shows addresses profile to domain mapping	649
Configure Interface	Configures MVR interface type and immediate leave mode; also displays MVR operational and active status	652
Configure Port	Configures MVR attributes for a port	652
Configure Trunk	Configures MVR attributes for a trunk	652
Configure Static Group Member		654
Add	Statically assigns MVR multicast streams to an interface	654
Show	Shows MVR multicast streams assigned to an interface	654
Show Member	Shows the multicast groups assigned to an MVR VLAN, the source address of the multicast services, and the interfaces with active subscribers	656
Show Statistics		657
Show Query Statistics	Shows statistics for query-related messages	657
Show VLAN Statistics	Shows statistics for protocol messages, number of active groups	657
Show Port Statistics	Shows statistics for protocol messages, number of active groups	657
Show Trunk Statistics	Shows statistics for protocol messages, number of active groups	657

* ECS4150-28P and ECS4150-54P only

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.
- ◆ [Configuring CPU Guard](#) – Sets thresholds in terms of CPU usage time and number of packets processed per second.
- ◆ [Displaying Memory Utilization](#) – Shows memory utilization parameters.
- ◆ [Resetting the System](#) – Restarts the switch immediately, at a specified time, after a specified delay, or at a periodic interval.

Displaying System Information

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

Parameters

These parameters are displayed:

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

Web Interface

To configure general system information:

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

Figure 2: System Information

The screenshot shows a web interface titled "System > General". It contains a table with the following fields and values:

System Description	ECS4150-28P
System Object ID	1.3.6.1.4.1.209.18.1.55.102
System Up Time	8 days, 8 hours, 15 minutes, and 36.17 seconds
System Name	
System Location	
System Contact	

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.

Management Software Information

- ◆ **Role** – Shows that this switch is operating as Master or Slave.
- ◆ **Loader Version** – Version number of loader code.
- ◆ **Linux Kernel Version** – Version number of Linux kernel.
- ◆ **Operation Code Version** – Version number of runtime code.

Temperature List

- ◆ **Thermal Detector** – The number of the detector on the main board.
- ◆ **Temperature** – Temperature at specified thermal detection point.

Web Interface

To view hardware and software version information.

1. Click System, then Switch.

Figure 3: General Switch Information



The screenshot shows a web interface titled 'System > Switch'. It contains two main sections: 'Main Board Information' and 'Management Software Information'. Below these is a 'Temperature List' table.

Main Board Information	
Serial Number	E022E3000614
Number of Ports	28
Hardware Version	R01
Main Power Status	VO

Management Software Information	
Role	Master
Loader Version	9.0.0.8
Linux Kernel Version	4.14.76-19.0.0
Operation Code Version	1.1.0.9

Temperature List Total: 1	
Thermal Detector	Temperature (°C)
1	31

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 4: 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 247.](#))
- ◆ **Static Entry Individual Port** – This switch allows static filtering for unicast and multicast addresses. (Refer to [“Setting Static Addresses” on page 205.](#))
- ◆ **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 161.](#))
- ◆ **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.

Web Interface

To view Bridge Extension information:

1. Click System, then Capability.

Figure 5: Displaying Bridge Extension Configuration

The screenshot shows the 'System > Capability' web interface. It has a left sidebar with 'General Capability' and 'Bridge Extension' (highlighted in blue). The main area displays configuration for 'Bridge Extension' with the following settings:

Setting	Value
Jumbo Frame	☒ Enabled
Extended Multicast Filtering Services	No
Traffic Classes	Enabled
Static Entry Individual Port	Yes
VLAN Version Number	2
VLAN Learning	T/C
Local VLAN Capable	No
Configurable PVID Tagging	Yes
Max Supported VLAN Numbers	4094
Max Supported VLAN ID	4094

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

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/FTPS/SFTP/SCP/TFTP or HTTP

Use the System > File (Copy) page to upload/download firmware or configuration settings using FTP, FTPS, SFTP, SCP, TFTP or HTTP. By backing up a file to a FTP/FTPS/SFTP/SCP/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/FTPS/SFTP 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.

- ◆ Secure Shell FTP (SFTP) provides a method of transferring files between two network devices over an SSH2-secured connection. SFTP functions similar to Secure Copy (SCP), using SSH for user authentication and data encryption.

Although the underlying premises of SFTP are similar to SCP, it requires some additional steps to verify the protocol versions and perform security checks. SFTP connection setup includes verification of the DSS signature, creation of session keys, creation of client-server and server-client ciphers, SSH key exchange, and user authentication. An SFTP channel is then opened, the SFTP protocol version compatibility verified, and SFTP finally initialized.

- ◆ 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.
 - FTPS Upload – Copies a file from an FTPS server to the switch.
 - FTPS Download – Copies a file from the switch to an FTPS 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.
 - SFTP Upload – Copies a file from an SFTP server to the switch.
 - SFTP Download – Copies a file from the switch to an SFTP server.
 - SCP Upload – Copies a file from an SCP server to the switch.
 - SCP Download – Copies a file from the switch to an SCP server.
 - 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/FTPS/SFTP/SCP/TFTP Server IP Address** – The IP address of an FTP/FTPS/SFTP/SCP/TFTP server.
- ◆ **User Name** – The user name for SFTP/SCP/FTP/FTPS server access.
- ◆ **Password** – The password for SFTP/SCP/FTP/FTPS 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, FTPS, SFTP, SCP, or TFTP Upload as the file transfer method.
4. If FTP, FTPS, SFTP, SCP, or TFTP Upload is used, enter the IP address of the file server.
5. If FTP/FTPS/SCP/SFTP Upload is used, enter the user name and password for your account on the FTP/FTPS/SCP/SFTP 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 6: Copy Firmware

The screenshot shows the 'System > File' web interface. At the top, there is a breadcrumb 'System > File'. Below it, an 'Action:' dropdown menu is set to 'Copy'. The main section is titled 'Copy Type' and contains several fields: 'Copy Type' (dropdown set to 'FTP Upload'), 'FTP Server IP Address' (text box with '192.168.1.123'), 'User Name' (text box with 'Edge01'), 'Password' (text box with masked characters), 'File Type' (dropdown set to 'Operation Code'), 'Source File Name' (text box with 'runtime-new.bin'), and 'Destination File Name' (radio button selected for 'runtime.bin' with a dropdown, and an empty text box with a radio button). At the bottom right, there are 'Apply' and 'Revert' buttons.

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

Saving the Running Configuration to a Local File

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

Parameters

The following parameters are displayed:

- ◆ **Copy Type** – The copy operation includes this option:
 - **Running-Config** – Copies the current configuration settings to a local file on the switch.
- ◆ **Destination File Name** – Copy to the currently designated startup file, or to a new file. The file name should not contain slashes (\ or /), the leading letter of the file name should not be a period (.), and the maximum length for file names is 32 characters. (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 7: Saving the Running Configuration

System > File

Actions: Copy

Copy Type: Running-Config

Destination File Name: startup1.cfg

Apply Revert

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 8: Setting Start-Up Files

System > File

Action: Set Start-Up

File List Total: 3

	File Name	File Type	Start-up	Modify Time	Size (bytes)
☒	runtime_v1.18.8.bin	Operation Code	Y	2023-02-07 14:40:05	51891848
☐	factory_default_config.cfg	Config File	N	2023-01-17 01:08:16	488
☒	startup1.cfg	Config File	Y	2023-02-17 07:39:01	1613

Apply Revert

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 9: Displaying System Files

System > File					
Action: Show					
File List Total: 3					
	File Name	File Type	Start-Up	Modify Time	Size (bytes)
☐	runtime_vr1.1.8.2.soc	Operation Code	Y	2023-02-07 14:40:05	51893946
☐	Factory_Default_Config.cfg	Config File	N	2023-01-17 01:39:19	455
☐	startup1.sfg	Config File	Y	2023-02-17 07:39:01	1673

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 *ECS4150.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 *ECS4150.BIX* from the server even though *ECS4150.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, *ecs4150.bix* and *ECS4150.bix* are considered to be unique files. Thus, if the upgrade file is stored as *ECS4150.bix* (or even *EcS4150.bix*) on a case-sensitive server, then the switch (requesting *ecs4150.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 *ECS4150.bix* filename must not be included since it is automatically appended by the switch. (Options: ftp, sftp, tftp)
- ◆ **Automatic Opcode Upgrade Reload** – Restarts the switch automatically after an opcode upgrade is completed.

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 10: Configuring Automatic Code Upgrade

The screenshot shows a web-based configuration interface. At the top, there is a breadcrumb 'System > File'. Below it, a section titled 'Automatic Operation Code Upgrade' contains the following settings:

- Automatic Opcode Upgrade:** A checkbox labeled 'Enabled' which is currently unchecked.
- Automatic Upgrade Location URL:** An empty text input field.
- Automatic Opcode Upgrade Reload:** A checkbox labeled 'Enabled' which is currently unchecked.

Below these settings is a note: 'Note: For automatic upgrades, the operation code file name must be set as EC34150.bin.' At the bottom right of the configuration area are two buttons: 'Apply' and 'Revert'.

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

```

.
.
Automatic Upgrade is looking for a new image
New image detected: current version 1.2.3.4; new version 1.2.3.5
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 11: Manually Setting the System Clock

The screenshot displays 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 shown as '2018-1-4 19:25:8'. The 'Maintain Type' is set to 'Manual' via a dropdown. The time fields are arranged in two rows: the first row contains 'Hours' (19), 'Minutes' (25), and 'Seconds' (8); the second row contains 'Month' (1), 'Day' (4), and 'Year' (2018). 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 12: Setting the Polling Interval for SNTP

The screenshot shows the 'System > Time' configuration page. At the top, there is a 'Steps' 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.

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 13: 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 'Current Time' is displayed as '2018-1-4 19:28:1'. The 'Maintain Type' is set to 'NTP' in a dropdown menu. Under the 'NTP Configuration' section, the 'Authentication Status' is shown with an unchecked checkbox labeled 'Enabled'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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

Specifying SNTP Time Servers

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

Parameters

The following parameters are displayed:

- ◆ **SNTP Server IP Address / Host Name** – Sets the IPv4 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 14: Specifying SNTP Time Servers

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** – Sets the IPv4 address for up to three 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 4)
- ◆ **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: 1-65533)

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 15: Adding an NTP Time Server




The screenshot shows the 'System > Time' configuration page. The 'Step' dropdown is set to '2. Configure Time Server' and the 'Action' dropdown is set to 'Add NTP Server'. The form contains three input fields: 'NTP Server IP Address/Host Name', 'Version' (with a value of '4' entered), and 'Authentication Key (1-65533)' (with an 'optional' label). 'Apply' and 'Revert' buttons are at the bottom right.

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 16: Showing the NTP Time Server List



The screenshot shows the 'System > Time' configuration page with the 'Action' dropdown set to 'Show NTP Server'. Below the form is a table titled 'NTP Server List Total: 1'.

	Server IP Address	Version	Authentication Key
☐	192.168.3.38	3	3

'Delete' and 'Revert' buttons are at the bottom right.

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: 1-65533)
- ◆ **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 17: Adding an NTP Authentication Key

The screenshot shows a web interface for configuring system time. The breadcrumb trail is 'System > Time'. The 'Step' dropdown is set to '2. Configure Time Server' and the 'Action' dropdown is set to 'Add NTP Authentication Key'. There are two input fields: 'Authentication Key (1-65533)' containing the number '3', and 'Key Context (1-32)' containing the string '51597N122103J066173M'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 18: 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 19: Setting the Time Zone

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 20: Configuring Summer Time

The screenshot shows a web-based configuration interface for a switch. The breadcrumb trail at the top is 'System > Time'. Below it, a step indicator shows 'Step: 4. Configure Summer Time'. The main configuration area includes:

- 'Summer Time in Effect' set to 'No'.
- 'Status' with an 'Enabled' checkbox.
- 'Name' with an empty text input field.
- 'Mode' with a dropdown menu showing 'Predefined'.
- 'Predefined Mode Configuration' section with 'Daylight Savings' set to 'Australia'.

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

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 21: Console Port Settings

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

- Login Timeout (10-300): 300 sec
- Exec Timeout (00-60000): ☒ 600 sec
- Password Threshold (1-120): ☒ 3
- Silent Time (1-65535): ☐ sec
- Data Bits: 8
- Stop Bits: 1
- Parity: None
- Speed: 115200 baud

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

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 22: Telnet Connection Settings

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

Parameter	Value	Unit
Telnet Status	☒ Enabled	
TCP Port (1-65535)	23	
Max Sessions (1-8)	8	
Login Timeout (10-300)	300	sec
Exec Timeout (60-65535)	600	sec
Password Threshold (1-120)	☒ 3	
Silent Time (1-65535)	☐	sec

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

Displaying CPU Utilization

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

Parameters

The following parameters are displayed:


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

Web Interface

To display CPU utilization:

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

Figure 23: Displaying CPU Utilization



Configuring CPU Guard

Use the System > CPU Guard page to set the CPU utilization high and low watermarks in percentage of CPU time utilized and the CPU high and low thresholds in the number of packets being processed per second.

Parameters

The following parameters are displayed:

- ◆ **CPU Guard Status** – Enables CPU Guard. (Default: Disabled)
- ◆ **High Watermark** – If the percentage of CPU usage time is higher than the high-watermark, the switch stops packet flow to the CPU (allowing it to catch up with packets already in the buffer) until usage time falls below the low watermark. (Range: 40-100 %; Default: 90 %)
- ◆ **Low Watermark** – If packet flow has been stopped after exceeding the high watermark, normal flow will be restored after usage falls beneath the low watermark. (Range: 40-100 %; Default: 70 %)
- ◆ **Maximum Threshold** – If the number of packets being processed by the CPU is higher than the maximum threshold, the switch stops packet flow to the CPU (allowing it to catch up with packets already in the buffer) until the number of packets being processed falls below the minimum threshold. (Range: 50-300 pps; Default: 300 pps)
- ◆ **Minimum Threshold** – If packet flow has been stopped after exceeding the maximum threshold, normal flow will be restored after usage falls beneath the minimum threshold. (Range: 50-300 pps; Default: 50 pps)

- ◆ **Trap Status** – If enabled, an alarm message will be generated when utilization exceeds the high watermark or exceeds the maximum threshold.
(Default: Disabled)

Once the high watermark is exceeded, utilization must drop beneath the low watermark before the alarm is terminated, and then exceed the high watermark again before another alarm is triggered.

Once the maximum threshold is exceeded, utilization must drop beneath the minimum threshold before the alarm is terminated, and then exceed the maximum threshold again before another alarm is triggered.

- ◆ **Current Threshold** – Shows the configured threshold in packets per second.

Web Interface

To configure CPU Guard:

1. Click System, CPU Guard.
2. Set CPU guard status, configure the watermarks or threshold parameter, enable traps if required.
3. Click Apply.

Figure 24: Configuring CPU Guard

The screenshot shows a web interface for configuring CPU Guard. The title bar indicates 'System > CPU Guard'. The configuration is as follows:

Parameter	Value	Unit
CPU Guard Status	☒ Enabled	
High Watermark (40-100)	80	%
Low Watermark (40-100)	70	%
Maximum Threshold (50-500)	300	packets/sec
Minimum Threshold (50-300)	50	packets/sec
Trap Status	☒ Enabled	
Current Threshold	300	packets/sec

Buttons: Apply, Revert

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	590,151,680 bytes	54%
Used Size	483,598,144 bytes	46%
Total	1,073,749,824 bytes	

Resetting the System

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

Command Usage

- ◆ This command resets the entire system.
- ◆ When the system is restarted, it will always run the Power-On Self-Test. It will also retain all configuration information stored in non-volatile memory. (See ["Saving the Running Configuration to a Local File" on page 81](#)).

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: 0-34560)
- **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 26: Restarting the Switch (Immediately)




Figure 27: Restarting the Switch (In)




Figure 28: 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 regularly time: 11:29 everyday.

System Reload Configuration:
Reset Mode:
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.

Figure 29: Restarting the Switch (Regularly)

System > Reset

System Reload Information:
No configured settings for reloading.

System Reload Configuration:
Reset Mode:
Time: 05:30 (HH:MM)
Period: ☒ Daily
☐ Weekly
☐ Monthly

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

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.
- ◆ [Displaying Statistics](#) – Shows Interface, Etherlike, and RMON port statistics in table or chart form.
- ◆ [Displaying Statistical History](#) – Displays statistical history for the specified interfaces.
- ◆ [Displaying Transceiver Data](#) – Displays identifying information, and operational parameters for optical transceivers which support DDM.
- ◆ [Configuring Transceiver Thresholds](#) – Configures thresholds for alarm and warning messages for optical transceivers which support DDM.
- ◆ [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.
- ◆ [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.
- ◆ [Flow Sampling](#) – Configures periodic sampling of traffic flows.
- ◆ [Traffic Segmentation](#) – Configures the uplinks and down links to a segmented group of ports.
- ◆ [Excluded VLAN](#) – Configures excluded VLAN sessions.
- ◆ [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

The Speed/Duplex mode is fixed at 1000full for Gigabit transceivers, and 10Gfull for 10 Gigabit transceivers.



Note: Auto-negotiation is not supported for SFP/SFP+/SFP28 transceivers.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28/54)
- ◆ **Type** – Indicates the port type. (1000BASE-T, 10GBASE SFP+, or 25GBASE SFP+)
- ◆ **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+/SFP28 ports.
 - **None** - Forced transceiver mode is not used for SFP/SFP+/SFP28 ports.
 - **SFP-Forced 1000SFP** - Always uses the SFP/SFP+/SFP28 port at 1000 Mbps, Full Duplex.
 - **SFP-Forced 10GSFP** - Always uses the SFP+/SFP28 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.

- **10f** - Supports 10 Mbps full-duplex operation.
- **100f** - Supports 100 Mbps full-duplex operation.
- **1000f** - Supports 1000 Mbps 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;
Advertised capabilities for
1000BASE-T – 10full, 100full, 1000full
1000BASE (SFP) – 1000full
10GBASE (SFP+) – 10Gfull

- ◆ **Speed/Duplex** – Displays the port speed and duplex mode. (100full, 10full)
- ◆ **Flow Control** – Enables flow control on the port. (Default: Disabled)

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 30: Configuring Connections by Port List

Interface > Port > General

Action: Configure by Port List

Port List Total: 34

Port	Type	Name	Admin	Media Type	Autonegotiation	Speed/Duplex	Flow Control
1	1000BASE-T		☒ Enabled	None	☒ Enabled ☒ 10f ☒ 1000f ☒ 100f ☐ Sym	100full	☐ Disabled
2	1000BASE-T		☒ Enabled	None	☒ Enabled ☒ 10f ☒ 1000f ☒ 100f ☐ Sym	100full	☐ Disabled
3	1000BASE-T		☒ Enabled	None	☒ Enabled ☒ 10f ☒ 1000f ☒ 100f ☐ Sym	100full	☐ Disabled

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

Parameters

Except for the trap command, refer to [“Configuring by Port List” on page 108](#) for more information on command usage and a description of the parameters.

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 31: Configuring Connections by Port Range

Interface > Port > General

Actions: Configure by Port Range

Port Range (1-54) -

Admin ☐ Enabled

Autonegotiation ☐ Enabled

☐ 10T ☐ 1000T ☐ FC

☐ 100T ☐ Sym

Speed Duplex 10Full

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 25GBASE SFP+)
- ◆ **Name** – Interface label.
- ◆ **Admin** – Shows if the port is enabled or disabled.

- ◆ **Oper Status** – Indicates if the link is Up or Down.
- ◆ **Shutdown Reason** – Shows the reason this interface has been shut down if applicable. Some of the reasons for shutting down an interface include being administratively disabled, or exceeding traffic boundary limits set by auto traffic control.
- ◆ **Media Type** – Shows the forced transceiver mode for SFP/SFP+/SFP28 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 32: Displaying Port Information

Port	Type	Name	Admin	Oper Status	Shutdown Reason	Media Type	Autonegotiation	Oper Speed Duplex	Oper Flow Control
1	1000BASE-T		Enabled	Up		None	Enabled	1000full	None
2	1000BASE-T		Enabled	Down	License	None	Enabled	1000full	None
3	1000BASE-T		Enabled	Down	License	None	Enabled	1000full	None
4	1000BASE-T		Enabled	Down	License	None	Enabled	1000full	None
5	1000BASE-T		Enabled	Down	License	None	Enabled	1000full	None
6	1000BASE-T		Enabled	Down	License	None	Enabled	1000full	None

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.
Received Unicast Packets	The number of subnetwork-unicast packets delivered to a higher-layer protocol.
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.
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 Octets	The total number of octets transmitted out of the interface, including framing characters.
Transmitted Errors	The number of outbound packets that could not be transmitted because of errors.
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.
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 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 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.
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.
<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.
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.
Internal MAC Transmit Errors	A count of frames for which transmission on a particular interface fails due to an internal MAC sublayer transmit error.
Pause Frames Input	If the counter is incrementing this means that the port is receiving pause frames.
Pause Frames Output	Pause outputs occur when the receiving port is getting overloaded and the device sends a pause request to the device connected to the port.
Symbol Errors	The number of times there was an invalid data symbol when a valid carrier was present.
<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 128-255 Byte Packets 256-511 Byte Packets 512-1023 Byte Packets 1024-1518 Byte Packets 1519-1536 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).	
CRC Align Errors		
<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 33: Showing Port Statistics (Table)



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 34: Showing Port Statistics (Chart)



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

Command Usage

- ◆ For a description of the statistics displayed on these pages, see [“Showing Port or Trunk Statistics” on page 111](#).
- ◆ To configure statistical history sampling, use the [“Displaying Statistical History” on page 116](#).

Parameters

These parameters are displayed:

Add

- ◆ **Port** – Port number. (Range: 1-28/54)

- ◆ **History Name** – Name of sample interval. (Range: 1-31 characters)
- ◆ **Interval** - The interval for sampling statistics. (Range: 1-86400 minutes)
- ◆ **Requested Buckets** - The number of samples to take. (Range: 1-96)

Show

- ◆ **Port** – Port number. (Range: 1-28/54)
- ◆ **History Name** – Name of sample interval. (Default settings: 15min, 1day)
- ◆ **Interval** - The interval for sampling statistics.
- ◆ **Requested Buckets** - The number of samples to take.

Show Details

- ◆ **Mode**
 - **Status** – Shows the sample parameters.
 - **Current Entry** – Shows current statistics for the specified port and named sample.
 - **Input Previous Entries** – Shows statistical history for ingress traffic.
 - **Output Previous Entries** – Shows statistical history for egress traffic.
- ◆ **Port** – Port number. (Range: 1-28/54)
- ◆ **Name** – Name of sample interval.

Web Configuration

To configure a periodic sample of statistics:

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

Figure 35: Configuring a History Sample

Interface > Port > History

Action: Add

Port: 1

History Name:

Interval (1-36480): seconds

Requested Buckets (1-60):

Apply Revert

To show the configured entries for a history sample:

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

Figure 36: Showing Entries for History Sampling

Interface > Port > History

Action: Show

Port: 1

History Name List Total: 3


	History Name	Interval	Requested Buckets
☐	10min	600	60
☐	1day	86400	7
☐	10P1	60	60

Delete Revert

To show the configured parameters for a sampling entry:

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


Figure 37: Showing Status of Statistical History Sample



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

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

Figure 38: Showing Current Statistics for a History Sample



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

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

Figure 39: Showing Ingress Statistics for a History Sample

Interface > Port > History

Action: [Show Details](#)

Mode: ☐ Status ☐ Current Entry ☒ Input Previous Entries ☐ Output Previous Entries

Port:

Name:

Input Previous Entry List Total: 10

Start Time	%	Octets	Unicast	Multicast	Broadcast	Errors
00d 00:00:01	0.00	523668	5774	837	90	0
00d 00:15:01	0.00	466711	5147	588	1	0
00d 00:30:01	0.00	425432	5138	590	3	0
00d 00:45:01	0.00	464422	5138	582	0	0
00d 01:00:01	0.00	425154	2908	513	0	0
00d 01:15:01	0.00	388886	2423	516	4	0
00d 01:30:01	0.00	354000	2475	588	1	0
00d 01:45:01	0.00	303167	2478	583	0	0
00d 02:00:01	0.00	367150	2560	587	0	0
00d 02:15:01	0.00	374615	2545	587	0	0

[Refresh](#)

Displaying Transceiver Data

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

Parameters

These parameters are displayed:

- ◆ **Port** – Port number. (Range: 1-28/54)
- ◆ **General** – Information on connector type and vendor-related parameters.
- ◆ **DDM Information** – Information on temperature, supply voltage, laser bias current, laser power, and received optical power.

The switch can display diagnostic information for SFP modules which support the SFF-8472 Specification for Diagnostic Monitoring Interface for Optical Transceivers. This information allows administrators to remotely diagnose


problems with optical devices. This feature, referred to as Digital Diagnostic Monitoring (DDM) provides information on transceiver parameters.

Web Interface

To display identifying information and functional parameters for optical transceivers:

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

Figure 40: Displaying Transceiver Data



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. (Range: 1-28/54)
- ◆ **General** – Information on connector type and vendor-related parameters.
- ◆ **DDM Information** – Information on temperature, supply voltage, laser bias current, laser power, and received optical power.

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

- ◆ **Trap** – Sends a trap when any of the transceiver's operation values falls outside of specified thresholds. (Default: Disabled)
- ◆ **Auto Mode** – Uses default threshold settings obtained from the transceiver to determine when an alarm or trap message should be sent. (Default: Enabled)
- ◆ **DDM Thresholds** – Information on alarm and warning thresholds. The switch can be configured to send a trap when the measured parameter falls outside of the specified thresholds.

The following alarm and warning parameters are supported:

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

The configurable ranges are:

- **Temperature:** -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 41: Configuring Transceiver Thresholds

	High Alarm	High Warning	Low Warning	Low Alarm
Temperature(°C)	75.00	70.00	0.00	-123.00
Voltage(Volts)	3.50	3.40	3.15	3.10
Current(mA)	100.00	90.00	7.00	6.00
Tx Power(dBm)	-9.00	-9.50	-11.50	-12.00
Rx Power(dBm)	-3.00	-3.50	-21.00	-21.50

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

Apply Revert

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 Time Domain Reflectometry (TDR). TDR detects a cable fault by sending a signal through the cable and reading the signal that is reflected back. TDR can only determine if a link is valid or faulty.

- ◆ Cable diagnostics can only be performed on twisted-pair media.
- ◆ This cable test is only accurate for Gigabit 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 those listed below. Note that TDR testing can only show Test failed or OK.
 - Test failed
 - OK: Correctly terminated pair
 - Open: Open pair, no link partner
 - Short: Shorted pair
 - Impedance mismatch: Terminating impedance is not in the reference range.
 - No cable
 - Not tested
 - 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.
 - Unknown – Unknown error
- ◆ 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)
- ◆ **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 42: Performing Cable Tests

Interface > Port > Cable Test

Cable Test Port List Total: 26

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	Not Tested	Not Tested	Not Tested	Not Tested		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 28 trunks at a time on the switch.

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


Command Usage

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

- ◆ Finish configuring trunks before you connect the corresponding network cables between switches to avoid creating a loop.
- ◆ You can create up to 28 trunks on a switch, 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 43: Configuring Static Trunks



Command Usage

- ◆ 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: 1-28)
- ◆ **Member** – The initial trunk member. Use the Add Member page to configure additional members.
 - **Unit** – Unit identifier. (Range: 1)
 - **Port** – Port identifier. (Range: 1-28/54)

Web Interface

To create a static trunk:

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

Figure 44: Creating Static Trunks



The screenshot shows a web interface for configuring static trunks. At the top, the breadcrumb path is 'Interface > Trunk > Static'. Below this, there are two dropdown menus: 'Step' set to '1. Configure Trunk' and 'Action' set to 'Add'. The main form area contains a 'Trunk ID (1-28)' text input field. Below it, the 'Member' section has 'Unit' and 'Port' dropdown menus, both currently set to '1'. At the bottom right of the form are 'Apply' and 'Reset' buttons.

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 45: Adding Static Trunks Members

Interface > Trunk > Static

Step: 1. Configure Trunk Action: Add Member

Trunk: 1

Member: Unit: 1 Port: 2

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

Figure 46: 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	Auto negotiation	Speed Duplex	Flow Control
1	10GBASE-T		☒ Enabled	☒ Enabled ☒ 10b ☒ 100b ☒ 1000b ☒ 10G	10Gbit	☒ 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 47: Showing Information for Static Trunks

Trunk	Type	Name	Admin	Oper Status	Shutdown Reason	Autoaggregation	Oper Speed Duplex	Oper Flow Control
1	100BASE-T	Enabled	Enabled	Down		Enabled	1000Kb	None

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 48: 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.
- ◆ 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 key is set to the same value as the operational key of the first member port (see the “show lacp internal” command in the *CLI Reference Guide*).

Parameters

These parameters are displayed:

Configure Aggregator

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

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

- ◆ **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.

- ◆ **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.
- ◆ **System MAC Address** – The device MAC address assigned to each trunk.

Configure Aggregation Port - General

- ◆ **Port** – Port identifier. (Range:1-28/54)
- ◆ **LACP Status** – Enables or disables LACP on a port.

Configure Aggregation Port - Actor/Partner

- ◆ **Port** – Port number. (Range: 1-28/54)
- ◆ **Mode** – Specifies the port's LACP activity mode.
 - **Active** – Specifies the port's activity mode to initiate and transmit LACP negotiation packets.
 - **Passive** – Specifies the port's activity mode to only respond to LACP negotiation packets.
- ◆ **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.

If the admin key is not set, the actor's operational key is determined by the port's link speed (10G - 5, 1000f - 4, 100f - 3, 10f - 2).

- ◆ **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 49: Configuring the LACP Aggregator Admin Key

Trunk	Admin Key (1-4096)	Timeout Mode	System Priority	System MAC Address
1		Long Timeout	32768	90-3C-83-8C-44-82
2		Long Timeout	32768	90-3C-83-8C-44-82
3		Long Timeout	32768	90-3C-83-8C-44-82
4		Long Timeout	32768	90-3C-83-8C-44-82
5		Long Timeout	32768	90-3C-83-8C-44-82

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 50: Enabling LACP on a Port

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 51: Configuring LACP Parameters on a Port

Port	Mode	Admin Key (0-4095)	System Priority (0-4095)	Port Priority (0-4095)
1	Active	4	32768	32768
2	Active	4	32768	32768
3	Active	4	32768	32768
4	Active	4	32768	32768
5	Active	4	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 52: Showing Members of a Dynamic Trunk

The screenshot shows the 'Interface > Trunk > Dynamic' configuration page. The 'Steps' dropdown is set to '3. Configure Trunk' and the 'Actions' dropdown is set to 'Show Member'. Below this, there is a 'Trunk' dropdown set to '1'. A 'Member List' section shows 'Total: 2' members. The table below lists the members:

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

Figure 53: Configuring Connection Settings for a Dynamic Trunk

The screenshot shows the 'Interface > Trunk > Dynamic' configuration page with the 'Steps' dropdown set to '3. Configure Trunk' and the 'Actions' dropdown set to 'Configure'. Below this, there is a 'Dynamic Trunk List' section showing 'Total: 1' trunk. The table below lists the trunk configuration:

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

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

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 54: Showing Connection Parameters for Dynamic Trunks



The screenshot shows a web-based configuration interface. At the top, there's a breadcrumb trail: 'Interface > Trunk > Dynamic'. Below this, there are two dropdown menus: 'Step: 3 Configure Trunk' and 'Action: Show'. A table titled 'Dynamic Trunk List Total: 1' displays the following data:

Trunk	Type	Name	Admin	Oper Status	Shutdown Reason	Autonegotiation	Oper Speed	Duplex	Oper Flow Control	Link Up/Down Trst
3	1000BASE-SFP		Enabled	Up		Enabled	1000full		None	Enabled

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

Parameters

These parameters are displayed:

Table 7: LACP Port Counters

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

Web Interface

To display LACP port counters:

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

5. Select a group member from the Port list.

Figure 55: 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

LACPDU's Sent	29	LACPDU's 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	The number of seconds between periodic LACPDU transmissions.
Admin State, Oper State	Administrative or operational values of the actor's state parameters: ◆ Expired – The actor's receive machine is in the expired state; ◆ Defaulted – The actor's receive machine is using defaulted operational partner information, administratively configured for the partner. ◆ Distributing – If false, distribution of outgoing frames on this link is disabled; i.e., distribution is currently disabled and is not expected to be enabled in the absence of administrative changes or changes in received protocol information. ◆ Collecting – Collection of incoming frames on this link is enabled; i.e., collection is currently enabled and is not expected to be disabled in the absence of administrative changes or changes in received protocol information. ◆ Synchronization – The System considers this link to be IN_SYNC; i.e., it has been allocated to the correct Link Aggregation Group, the group has been associated with a compatible Aggregator, and the identity of the Link Aggregation Group is consistent with the System ID and operational Key information transmitted.

Table 8: LACP Internal Configuration Information (Continued)

Parameter	Description
Admin State, Oper State (continued)	◆ Aggregation – The system considers this link to be aggregatable; i.e., a potential candidate for aggregation. ◆ Long timeout – Periodic transmission of LACPDU's 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 56: Displaying LACP Port Internal Information

The screenshot shows a web interface for configuring LACP. The breadcrumb navigation is 'Interface > Trunk > Dynamic'. Below this, there are two dropdown menus: 'Steps' set to '2. Configure Aggregation Port' and 'Action' set to 'Show Information'. There are three radio buttons: 'Counters' (unselected), 'Internal' (selected), and 'Neighbors' (unselected). Below the radio buttons, there is a 'Port' dropdown menu set to '1' and a 'Trunk ID' field set to '2'. The main section is titled 'Port Internal Information' and contains a table of LACP parameters and their values.

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 57: Displaying LACP Port Remote Information

Interface > Trunk > Dynamic

Step: 2. Configure Aggregation Port Actions: 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 58: Configuring Load Balancing



Saving Power

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

Command Usage

◆ 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 59: Enabling Power Savings



Configuring Local Port Mirroring

Use the Interface > 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 60: Configuring Local Port Mirroring



Command Usage

- ◆ The switch supports up to 6 port mirror sessions.
- ◆ 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 144](#)).

- ◆ Monitor port speed should match or exceed source port speed, otherwise traffic may be dropped from the monitor port.
- ◆ 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, 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 61: Configuring Local Port Mirroring

The screenshot shows a web interface for configuring port mirroring. The title bar reads 'Interface > Mirror'. Below the title bar is an 'Actions' dropdown menu with 'Add' selected. The main configuration area contains three rows of fields: 'Source Port' with 'Unit 1' and 'Port 1' dropdowns, 'Target Port' with 'Unit 1' and 'Port 1' dropdowns, and 'Type' with a 'Both' dropdown. At the bottom right of the form are two buttons: 'Apply' and 'Reset'.

To display the configured mirror sessions:

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


Figure 62: Displaying Local Port Mirror Sessions



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 63: 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 142](#)), 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 165](#)) 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.

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 167](#)).

- MAC address learning is not supported on RSPAN uplink ports when RSPAN is enabled on the switch. Therefore, even if spanning tree is enabled after RSPAN has been configured, MAC address learning will still not be re-started on the RSPAN uplink ports.
- *IEEE 802.1X* – RSPAN and 802.1X are mutually exclusive functions. When 802.1X is enabled globally, RSPAN uplink ports cannot be configured, even though RSPAN source and destination ports can still be configured. When RSPAN uplink ports are enabled on the switch, 802.1X cannot be enabled globally.
- *Port Security* – If port security is enabled on any port, that port cannot be set as an RSPAN uplink port, even though it can still be configured as an RSPAN source or destination port. Also, when a port is configured as an RSPAN uplink port, port security cannot be enabled on that port.

Parameters

These parameters are displayed:

- ◆ **Session** – A number identifying this RSPAN session. (Range: 1)
Only one active session is allowed.
- ◆ **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 165](#)).
- ◆ **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 64: Configuring Remote Port Mirroring (Source)

The screenshot shows the 'Interface > RSPAN' configuration page. The 'Switch Role' is set to 'Source'. Below the configuration fields is a table titled 'Source Port Configuration List' with a total of 25 rows. The table has two columns: 'Source Port' and 'Type'. The first four rows are visible, showing ports 1, 2, 3, and 4, all with 'None' selected for the Type.

Source Port	Type
1	None
2	None
3	None
4	None

Figure 65: Configuring Remote Port Mirroring (Intermediate)

The screenshot shows the 'Interface > RSPAN' configuration page. The 'Session' is set to 1, 'Operation Status' is Down, 'Switch Role' is Intermediate, and 'Remote VLAN' is 9. Below the configuration fields, there is a table titled 'Uplink Port List' with a total of 26 ports. The table has two columns: 'Port' and 'Uplink'. The 'Port' column lists ports 1, 2, 3, and 4. The 'Uplink' column has checkboxes for each port, all of which are currently unchecked.

Port	Uplink
1	☐
2	☐
3	☐
4	☐

Figure 66: Configuring Remote Port Mirroring (Destination)

The screenshot shows the 'Interface > RSPAN' configuration page. The 'Session' is set to 1, 'Operation Status' is Down, 'Switch Role' is Destination, 'Destination Port' is 1, 'Tag' is Untagged, and 'Remote VLAN' is 9. Below the configuration fields, there is a table titled 'Uplink Port List' with a total of 26 ports. The table has two columns: 'Port' and 'Uplink'. The 'Port' column lists ports 1, 2, 3, and 4. The 'Uplink' column has checkboxes for each port, all of which are currently unchecked.

Port	Uplink
1	☐
2	☐
3	☐
4	☐

Sampling Traffic Flows

The flow sampling (sFlow) feature embedded on this switch, together with a remote sFlow Collector, can provide network administrators with an accurate, detailed and real-time overview of the types and levels of traffic present on their network. The sFlow Agent samples 1 out of n packets from all data traversing the switch, re-encapsulates the samples as sFlow datagrams and transmits them to the sFlow Collector. This sampling occurs at the internal hardware level where all traffic is seen, whereas traditional probes will only have a partial view of traffic as it is sampled at the monitored interface. Moreover, the processor and memory load imposed by the sFlow agent is minimal since local analysis does not take place. The wire-speed transmission characteristic of the switch is thus preserved even at high traffic levels.



Note: The terms “collector”, “receiver” and “owner”, in the context of this chapter, all refer to a remote server capable of receiving the sFlow datagrams generated by the sFlow agent of the switch.

As the Collector receives streams from the various sFlow agents (other switches or routers) throughout the network, a timely, network-wide picture of utilization and

traffic flows is created. Analysis of the sFlow stream(s) can reveal trends and information that can be leveraged in the following ways:

- ◆ Detecting, diagnosing, and fixing network problems
- ◆ Real-time congestion management
- ◆ Understanding application mix (P2P, Web, DNS, etc.) and changes
- ◆ Identification and tracing of unauthorized network activity
- ◆ Usage accounting
- ◆ Trending and capacity planning

Configuring sFlow Receiver Settings Use the Interface > sFlow (Configure Receiver – Add) page to create an sFlow receiver on the switch.

Parameters

These parameters are displayed:

- ◆ **Receiver Owner Name**² – The name of the receiver. (Range: 1-256 characters; Default: None)
- ◆ **Receiver Timeout** – The time that the sFlow process will continuously send samples to the Collector before resetting all sFlow port parameters. (Range: 30-10000000 seconds, where 0 indicates no time out)

The sFlow parameters affected by this command include the sampling interval, the receiver's name, address and UDP port, the time out, maximum header size, and maximum datagram size.
- ◆ **Receiver Destination**² – IP address of the sFlow Collector.
 - *ipv4-address* - IPv4 address of the sFlow collector. Valid IPv4 addresses consist of four decimal numbers, 0 to 255, separated by periods.
 - *ipv6-address* - IPv6 address of the sFlow collector. A full IPv6 address including the network prefix and host address bits. An IPv6 address consists of 8 colon-separated 16-bit hexadecimal values. One double colon may be used to indicate the appropriate number of zeros required to fill the undefined fields.
- ◆ **Receiver Socket Port**² – The UDP port on which the sFlow Collector is listening for sFlow streams. (Range: 1-65535)
- ◆ **Maximum Datagram Size** – Maximum size of the sFlow datagram payload. (Range: 200-1500 bytes)

2. Sampling must be disabled by setting the time out to 0 before these fields can be configured.

- ◆ **Datagram Version** – Sends either v4 or v5 sFlow datagrams to the receiver.

Web Interface

To configure an sFlow receiver:

1. Click Interface, sFlow.
2. Select Configure Receiver from the Step list.
3. Select Add from the Action list.
4. Fill in the parameters for the sFlow receiver and monitored traffic.
5. Click Apply.

Figure 67: Configuring an sFlow Receiver

Interface > sFlow

Step: 1. Configure Receiver Actions: Add

Receiver Owner Name: sflow_server1

Receiver Timeout (30 - 18000000): 100 sec

Receiver Destination: 192.168.220.225

Receiver Socket Port (1 - 65535): 20500

Maximum Datagram Size (200 - 1500): 512 bytes

Datagram Version: ☒ v5

Apply Revert

Web Interface

To show configured receivers:

1. Click Interface, sFlow.
2. Select Configure Receiver from the Step list.
3. Select Show from the Action list.

Figure 68: Showing sFlow Receivers

Interface > sFlow

Step: 1. Configure Receiver Actions: Show

Receiver List Time: 1

	Owner Name	Timeout	Destination	Socket Port	Maximum Datagram Size	Datagram Version
	sflow_server1	67	192.168.220.225	20500	512	5

Details Revert

Configuring an sFlow Polling Instance Use the Interface > sFlow (Configure Details – Add) page to enable an sFlow polling data source that polls periodically based on a specified time interval, or an sFlow data source instance that takes samples periodically based on the number of packets processed.

Parameters

These parameters are displayed in the web interface:

- ◆ **Receiver Owner Name** – The name of the receiver. (Range: 1-256 characters; Default: None)
- ◆ **Type** – Specifies the polling type as an sFlow polling data source for a specified interface that polls periodically based on a specified time interval, or an sFlow data source instance for a specific interface that takes samples periodically based on the number of packets processed.
- ◆ **Data Source** – The source from which the samples will be taken and sent to a collector.
- ◆ **Instance ID** – An instance ID used to identify the sampling source. (Range: 1)
- ◆ **Sampling Rate** – The number of packets out of which one sample will be taken. (Range: 256-16777215 packets; Default: Disabled)
- ◆ **Maximum Header Size** – Maximum size of the sFlow datagram header. (Range: 64-256 bytes)

Web Interface

To configure an sFlow sampling or polling instance:

1. Click Interface, sFlow.
2. Select Configure Details from the Step list.
3. Select Add from the Action list.
4. Fill in the parameters for the sFlow instance, including sampling rate and maximum header size.
5. Click Apply.

Figure 69: Configuring an sFlow Instance

The screenshot shows the 'Interface > sFlow' configuration page. The 'Step' dropdown is set to '2. Configure Details' and the 'Action' dropdown is set to 'Add'. The configuration fields are as follows:

- Receiver Owner Name:
- Type: ☒ Sampling ☐ Polling
- Data Source: Unit: Port:
- Instance ID (1-1):
- Sampling Rate (256-16777215):
- Maximum Header Size (64-256): bytes

Buttons: Apply, Revert

Web Interface

To show configured instances:

1. Click Interface, sFlow.
2. Select Configure Details from the Step list.
3. Select Show from the Action list.
4. Select the owner name from the scroll-down list.
5. Select sFlow type as Sampling or Polling.

Figure 70: Showing sFlow Instances

The screenshot shows the 'Interface > sFlow' configuration page with the 'Step' dropdown set to '2. Configure Details' and the 'Action' dropdown set to 'Show'. The configuration fields are as follows:

- Receiver Owner Name:
- Type: ☒ Sampling ☐ Polling

SamplingList Total: 1

	Data Source (Unit/Port)	Instance ID	Rate	Maximum Header Size (bytes)
☐	1/1	1	256	256

Buttons: Delete, Revert

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 71: Enabling Traffic Segmentation

Configuring Uplink and Downlink Ports

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

Command Usage

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

Table 10: Traffic Segmentation Forwarding

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

* The forwarding state for uplink-to-uplink ports is configured on the Configure Global page (see [page 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: 1-28/54)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-28)

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 72: Configuring Members for Traffic Segmentation

The screenshot shows a web interface titled "Interface > Traffic Segmentation". At the top, there is a "Step:" dropdown menu set to "2. Configure Session" and an "Action:" dropdown menu set to "Add". Below this, the "Session ID (1-4)" field is empty. The "Direction" dropdown menu is set to "Uplink". Under the "Interface" section, the "Port (1-28)" radio button is selected, and the "Trunk (1-28)" radio button is unselected. Both the "Port" and "Trunk" sections have empty input fields for specifying the identifier. At the bottom right, there are "Apply" and "Revert" buttons.

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 73: Showing Traffic Segmentation Members



Excluded VLAN

Excluded VLANs provide port-based security and isolation between ports within an assigned session. An Excluded VLAN session contains Uplink ports that can communicate with all other ports in the session, and Downlink ports that can only communicate with Uplink ports in the session. The Uplink ports are intended to provide open access to an external network, such as the Internet, while the Downlink ports provide restricted access to local users.

Command Usage

- ◆ An Excluded VLAN session consists of defined Downlink ports, Uplink ports, and VLANs. Up to 4 Excluded VLAN sessions can be configured on the switch, and up to 8 VLANs can be included in each session.
- ◆ Packets from a Downlink port can only be forwarded to Uplink ports in the same session.
- ◆ Packets from Uplink ports can be forwarded to any ports in the same session.
- ◆ Uplink and Downlink interfaces can include configured port trunks.
- ◆ A specified Uplink or Downlink interface can only be assigned to one Excluded VLAN session, it cannot be configured for multiple sessions at the same time.

Parameters

These parameters are displayed:

- ◆ **Session ID** – Excluded VLAN session. (Range: 1-4)
- ◆ **Direction** – Add an interface to the session by setting the direction to Uplink or Downlink. (Default: Uplink)
- ◆ **VLAN** – Specifies a VLAN ID. (Range: 1-4094)
- ◆ **VLAN Mask** – Specifies a binary bitmask that is applied to the VLAN ID to define a range of VLANs. When a bit of the VLAN Mask is 1, the value of the corresponding bit of the VLAN ID remains the same. When a bit of configured VLAN Mask is 0, the value of the corresponding bit of the VLAN ID is ignored. For example, a VLAN ID of 1 (000000000001) with a VLAN Mask of 4092 (111111111100) defines a range of VLAN IDs of 1-3. A VLAN Mask of 4095 (111111111111) defines a single VLAN ID. (Range: 0-4095)
- ◆ **Interface** – Specifies a range of ports or trunks.
 - **Port** – Port identifier or range. (Range: 1-28/54)
 - **Trunk** – Trunk identifier or range. (Range: 1-28)

Web Interface

To enable an excluded VLAN session:

1. Click Interface, Exclude VLAN.
2. Select Add from the Action list.
3. Configure a Session ID.
4. Set the interface direction, either Uplink or Downlink.
5. Set a VLAN ID or a VLAN ID range.
6. Specify the port or trunk interfaces.
7. Click Apply.

Figure 74: Configuring Excluded VLANs

Interface > Exclude Vlan

Action: Add

Session ID (1-4):

Direction: Uplink

VLAN (1-4094):

VLAN Mask (0-4095):

Interface: ☒ Port (1-12) - ☐ Trunk (1-16) -

Apply Revert

To display the configured excluded VLAN sessions:

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

Figure 75: Displaying Excluded VLANs

Interface > Exclude Vlan

Action: Show

Session List Total: 1

	Session ID	Direction	Interface	Vlan
☐	1	Uplink	Unit 1 / Port 0	23/4095

Delete Revert

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 76: 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 167](#)). 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: 1-28/54)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-28)
- ◆ **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 77: Configuring VLAN Trunking

Interface > VLAN Trunking

Interface: ☒ Port ☐ Trunk

Port VLAN Trunking List Total: 26

Port	VLAN Trunking Status
1	☐ Enabled
2	☐ Enabled
3	☐ Enabled
6	☐ Enabled
7	☐ Enabled
8	☐ Enabled
9	☐ Enabled
10	☐ Enabled
11	☐ Enabled
12	☐ 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.
- ◆ [L2PT Tunneling](#) – Configures Layer 2 Protocol Tunneling for the specified protocol.
- ◆ [Protocol VLANs](#)³ – Configures VLAN groups based on specified protocols.
- ◆ [IP Subnet VLANs](#)³ – Maps untagged ingress frames to a specified VLAN if the source address is found in the IP subnet-to-VLAN mapping table.
- ◆ [MAC-based VLANs](#)³ – 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.

3. 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.

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

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

This switch supports the following VLAN features:

- ◆ Up to 4094 VLANs based on the IEEE 802.1Q standard
- ◆ Distributed VLAN learning across multiple switches using explicit or implicit tagging 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 78: 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 167](#)). But you can still enable GVRP on these edge switches, as well as on the core switches in the network.

Figure 79: 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 144](#)). 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.
- ◆ **Status** – Enables or disables the specified VLAN.
- ◆ **Remote VLAN** – Reserves this VLAN for RSPAN (see [“Configuring Remote Port Mirroring” on page 144](#)).

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 673](#)).

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 144](#)).
- ◆ **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. Check Status to configure the VLAN as operational.
5. Specify whether the VLANs are to be used for remote port mirroring.
6. Click Apply.

Figure 80: Creating Static VLANs

The screenshot shows the 'VLAN > Static' configuration window. 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 and labeled 'Enabled'. 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 81: Modifying Settings for Static VLANs

The screenshot shows the 'VLAN > Static' configuration window. The 'Action' dropdown is set to 'Modify'. The 'VLAN ID (1-4094)' field has a dropdown arrow and shows '1'. 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 82: Showing Static VLANs



The screenshot shows a web interface for configuring static VLANs. At the top, there's a breadcrumb 'VLAN > Static' and an 'Action' dropdown menu set to 'Show'. Below this, a table titled 'Static VLAN List: Total 2' displays the following data:

VLAN ID	VLAN Name	Status	Access VLAN	LS Interface
1	DefaultVlan	Enabled	Disabled	Enabled
2	RAD	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: 1-28/54)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-28)
- ◆ **Mode** – Indicates VLAN membership mode for an interface. (Default: Hybrid)
 - **Access** - Sets the port to operate as an untagged interface. The port transmits and receives untagged frames on a single VLAN only.
 - **Hybrid** – Specifies a hybrid VLAN interface. The port may transmit tagged or untagged frames.
 - **1Q Trunk** – Specifies a port as an end-point for a VLAN trunk. A trunk is a direct link between two switches, so the port transmits tagged frames that

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

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

When using Access mode, and an interface is assigned to a new VLAN, its PVID is automatically set to the identifier for that VLAN.

- ◆ **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 port default VLAN if not matched to a configured MAC VLAN, IP-subnet VLAN, or protocol 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: Enabled)
 - If ingress filtering is disabled and a port receives frames classified to VLANs for which it is not a member, these frames will be flooded to all other ports that are members of the VLANs.
 - 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 cannot be included as a member of the VLAN via GVRP. For more information, see [“Configuring Dynamic VLAN Registration” on page 171](#).

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: 1-28/54)
- ◆ **Trunk Range** – Displays a list of ports. (Range: 1-28)



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 83: Configuring Static Members by VLAN Index

VLAN > Static

Action: Edit Member by VLAN

VLAN: 1

Interface: ☒ Port ☐ Trunk

Static VLAN Port Member List Total: 26

Port	Mode	PVID (1-4094)	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	☐	☒	☐	☐
6	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 84: Configuring Static VLAN Members by Interface

VLAN > Static

Action: Edit Member by Interface

Interface: ☒ Port ☐ Trunk

Mode: Hybrid

PVID: 1

Acceptable Frame Type: All

Ingress Filtering: ☒ Enabled

Static VLAN Membership List Total: 2

VLAN	Membership Type			
	Tagged	Untagged	Forbidden	None
1	☒	☐	☐	☐
6	☐	☒	☐	☐

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 85: Configuring Static VLAN Members by Interface Range

VLAN > Static

Action: Edit Member by Interface Range

Interface: ☒ Port ☐ Trunk

Port Range (1-28): -

Mode: Hybrid

VLAN ID (1-4094): -

Membership Type: ☒ Tagged ☐ Untagged ☐ Forbidden ☐ None

Apply Revert

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: 1-28/54)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-28)
- ◆ **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 167](#)).

- ◆ **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 86: Configuring Global Status of GVRP

VLAN > Dynamic

Step: 1. Configure General

GVRP Status ☒ Enabled

Apply Revert

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 87: Configuring GVRP for an Interface

VLAN > Dynamic

Step: 2. Configure Interface

Interface ☒ Port ☐ Trunk

Port List Total: 28

Port	GVRP Status	GVRP Timer (milliseconds)		
		Join (20-1000)	Leave (60-3000)	LeaveAll (200-10000)
1	☐ Enabled	20	60	1000
2	☐ Enabled	20	60	1000
3	☐ Enabled	20	60	1000
4	☐ 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 88: Showing Dynamic VLANs Registered on the Switch



The screenshot shows the 'VLAN > Dynamic' configuration page. The 'Steps' dropdown is set to '3. Show Dynamic VLAN' and the 'Actions' dropdown is set to 'Show VLAN'. Below the navigation bar, there is a table titled 'Dynamic VLAN List - Total: 2'. The table has three columns: 'VLANID', 'VLAN Name', and 'Status'.

VLANID	VLAN Name	Status
120		Enabled
155		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 89: Showing the Members of a Dynamic VLAN



The screenshot shows the 'VLAN > Dynamic' configuration page. The 'Steps' dropdown is set to '3. Show Dynamic VLAN' and the 'Actions' dropdown is set to 'Show VLAN Member'. Below the navigation bar, there is a dropdown menu for 'VLAN' with '120' selected. Below the dropdown, there is a table titled 'Dynamic VLAN Member List - Total: 10'. The table has one column: 'Interface'.

Interface
Unit 1 / Port 1
Unit 1 / Port 4
Unit 1 / Port 10
Unit 1 / Port 15
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 90: 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 180](#). 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 179](#).)
2. Create a Service Provider VLAN, also referred to as an SPVLAN (see [“Configuring VLAN Groups” on page 165](#)).
3. Configure the QinQ tunnel access port to Access mode (see [“Adding an Interface to a QinQ Tunnel” on page 182](#)).
4. Configure the QinQ tunnel access port to join the SPVLAN as an untagged member (see [“Adding Static Members to VLANs” on page 167](#)).
5. Configure the SPVLAN ID as the native VID on the QinQ tunnel access port (see [“Adding Static Members to VLANs” on page 167](#)).
6. Configure the QinQ tunnel uplink port to Uplink mode (see [“Adding an Interface to a QinQ Tunnel” on page 182](#)).

7. Configure the QinQ tunnel uplink port to join the SPVLAN as a tagged member (see [“Adding Static Members to VLANs” on page 167](#)).

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 182](#)). 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 etherypes 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 91: Enabling QinQ Tunneling

The screenshot shows a web-based configuration interface for VLAN > Tunnel. The title bar indicates the current step is '1. Configure Global'. Below this, there are two main configuration fields: 'Tunnel Status' with a checked 'Enabled' checkbox, and 'Ethernet Type' with a dropdown menu set to '100'. A note below the Ethernet Type field states '(800-FFFF, hexadecimal value)'. At the bottom right, there are two buttons: 'Apply' and 'Reset'.

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 175](#).
- ◆ 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:

- ◆ **Interface** – Port or trunk identifier.
- ◆ **C-VID (Customer VLAN ID)** – VLAN ID for the inner VLAN tag. (Range: 1-4094)
- ◆ **Assign S-VID (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 92: Configuring CVLAN to SPVLAN Mapping Entries

VLAN > Tunnel

Step: 2. Configure Service Action: Add

Interface: ☒ Unit 1 Port 2 ☐ Trunk 1

Match Type: ☒ Tagged

CVID (1-4094):

Action: Assign SVID (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 93: Showing CVLAN to SPVLAN Mapping Entries

VLAN > Tunnel

Step: 2. Configure Service Action: Show

Tunnel Service Subscriptions List Total: 1

	Interface	Customer VLAN ID	Service VLAN ID
☒	Eth 1/5	55	44

Default Service Total: 0

Discard Untagged Traffic Total: 0

Deletes Revert

The preceding example sets the SVID to 44 in the outer tag for egress packets exiting port 6 when the packet's CVID is 55. For a more detailed example, see the "switchport dot1q-tunnel service match cvlid" 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 179](#)). 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: 1-28/54)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-28)
- ◆ **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.
- ◆ **Priority Mapping** – Copies the inner tag priority to the outer tag priority. 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.


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 94: Adding an Interface to a QinQ Tunnel



L2PT Tunneling

When Layer 2 Protocol Tunneling (L2PT) is not used, protocol packets (e.g., STP) are flooded to 802.1Q access ports on the same edge switch, but filtered from 802.1Q tunnel ports. This creates disconnected protocol domains in the customer's network.

L2PT can be used to pass various types of protocol packets belonging to the same customer transparently across a service provider's network. In this way, normally segregated network segments can be configured to function inside a common protocol domain.

Command Usage

- ◆ L2PT encapsulates protocol packets entering ingress ports on the service provider's edge switch, replacing the destination MAC address with a proprietary MAC address (for example, the spanning tree protocol uses 10-12-CF-00-00-02), a reserved address for other specified protocol types (as defined in IEEE 802.1ad – Provider Bridges), or a user-defined address. All intermediate switches carrying this traffic across the service provider's network treat these encapsulated packets in the same way as normal data, forwarding them to the tunnel's egress port. The egress port decapsulates these packets, restores the proper protocol and MAC address information, and then floods them onto the same VLANs at the customer's remote site (via all of the appropriate tunnel ports and access ports⁴ connected to the same metro VLAN).
- ◆ The way in which L2PT processes packets is based on the following criteria – (1) packet is received on a QinQ uplink port, (2) packet is received on a QinQ

4. Access ports in this context are 802.1Q trunk ports.

access port, or (3) received packet is Cisco-compatible L2PT (i.e., as indicated by a proprietary MAC address).

Processing protocol packets defined in IEEE 802.1ad – Provider Bridges

- ◆ When an IEEE 802.1ad protocol packet is received on an uplink port (i.e., an 802.1Q tunnel ingress port connecting the edge switch to the service provider network)
 - with the destination address 01-80-C2-00-00-00,0B~0F (C-VLAN tag), it is forwarded to all QinQ uplink ports and QinQ access ports in the same S-VLAN for which L2PT is enabled for that protocol.
 - with the destination address 01-80-C2-00-00-01~0A (S-VLAN tag), it is filtered, decapsulated, and processed locally by the switch if the protocol is supported.
- ◆ When a protocol packet is received on an access port (i.e., an 802.1Q trunk port connecting the edge switch to the local customer network)
 - with the destination address 01-80-C2-00-00-00,0B~0F (C-VLAN), and
 - L2PT is enabled on the port, the frame is forwarded to all QinQ uplink ports and QinQ access ports on which L2PT is enabled for that protocol in the same S-VLAN.
 - L2PT is disabled on the port, the frame is decapsulated and processed locally by the switch if the protocol is supported.
 - with destination address 01-80-C2-00-00-01~0A (S-VLAN), the frame is filtered, decapsulated, and processed locally by the switch if the protocol is supported.

Processing Cisco-compatible protocol packets

- ◆ When a Cisco-compatible L2PT packet is received on an uplink port, and
 - recognized as a CDP/VTP/STP/PVST+ protocol packet (where STP means STP/RSTP/MSTP), it is forwarded to the following ports in the same S-VLAN: (a) all access ports for which L2PT has been disabled, and (b) all uplink ports.
 - recognized as a Generic Bridge PDU Tunneling (GBPT) protocol packet (i.e., having the destination address 01-00-0C-CD-CD-D0), it is forwarded to the following ports in the same S-VLAN:
 - other access ports for which L2PT is enabled after decapsulating the packet and restoring the proper protocol and MAC address information.
 - all uplink ports.
- ◆ When a Cisco-compatible L2PT packet is received on an access port, and
 - recognized as a CDP/VTP/STP/PVST+ protocol packet, and

- L2PT is enabled on this port, it is forwarded to the following ports in the same S-VLAN: (a) other access ports for which L2PT is enabled, and (b) uplink ports after rewriting the destination address to make it a GBPT protocol packet (i.e., setting the destination address to 01-00-0C-CD-CD-D0).
- L2PT is disabled on this port, it is forwarded to the following ports in the same S-VLAN: (a) other access ports for which L2PT is disabled, and (b) all uplink ports.
- recognized as a GBPT protocol packet (i.e., having the destination address 01-00-0C-CD-CD-D0), and
 - L2PT is enabled on this port, it is forwarded to other access ports in the same S-VLAN for which L2PT is enabled
 - L2PT is disabled on this port, it is forwarded to the following ports in the same S-VLAN: (a) other access ports for which L2PT is disabled, and (b) all uplink ports.

L2PT Functional Requirements

- ◆ For L2PT to function properly, QinQ must be enabled on the switch (see [“Enabling QinQ Tunneling on the Switch”](#)), and the interface configured to 802.1Q uplink mode (see [“Adding an Interface to a QinQ Tunnel”](#)).

Configuring the L2PT Tunnel Address

Use the VLAN > L2PT (Configure Global) page to configure the destination MAC address for Layer 2 Protocol Tunneling.

Parameters

These parameters are displayed:

Tunnel MAC Address – The switch rewrites the destination MAC address in all upstream L2PT protocol packets (e.g., STP BPDUs) to this value, and forwards them on to uplink ports. (Format xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx; Default: 01-12-CF-.00-00-02, proprietary tunnel address)

Web Interface

To configure the destination MAC address for L2PT:

1. Click VLAN, L2PT.
2. Select Configure Global from the Step list.
3. Set the tunnel MAC address.
4. Click Apply.

Figure 95: Configuring the L2PT Tunnel Address

The screenshot shows a web interface for configuring L2PT. The breadcrumb is 'VLAN > L2PT'. Below it, a 'Step:' dropdown menu is set to '1. Configure Global'. The main configuration area has a label 'Tunnel MAC Address' followed by a text input field containing '01-12-CF-00-00-02'. To the right of the input field is a placeholder text '(XX-XX-XX-XX-XX-XX 01 XXXXXXXXXXXX)'. At the bottom right of the form are two buttons: 'Apply' and 'Revert'.

Enabling L2PT for Selected Interfaces Use the VLAN > L2PT (Configure Interface) page to enable Layer 2 Protocol Tunneling on selected interfaces.

Parameters

These parameters are displayed:

- ◆ **Interface** – Port or trunk identifier.
- ◆ **CDP** - Cisco Discovery Protocol
- ◆ **LACP** - Link Aggreation Control Protocol
- ◆ **LLDP** - Link Layer Discovery Protocol
- ◆ **PVST+** - Cisco Per VLAN Spanning Tree Plus
- ◆ **Spanning Tree** - Spanning Tree (STP, RSTP, MSTP)
- ◆ **VTP** - Cisco VLAN Trunking Protocol

Web Interface

To enable the L2PT on selected interfaces:

1. Click VLAN, L2PT.
2. Select Configure Interface from the Step list.
3. Mark the protocols to be passed through the L2PT tunnel on the required interfaces.
4. Click Apply.

Figure 96: Enabling L2PT on Required Interfaces

Interface	CDP	LACP	LLDP	PVST+	Spanning Tree	VTP
Eth1/1	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled
Eth1/2	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled
Eth1/3	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled
Eth1/4	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled
Eth1/5	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled

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 165](#)). 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, PPPOE-DIS, PPPOE-SES, 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 97: 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 98: 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	86 35	2
☐	RFC 1042	08 06	1
☐	RFC 1042	86 35	3
☐	LLC 0801	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 167), 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: 1-28/54)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-28)
- ◆ **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 99: Assigning Interfaces to Protocol VLANs

The screenshot shows the 'VLAN > Protocol' configuration page. The 'Step' is '2. Configure Interface' and the 'Action' is 'Add'. The 'Interface' section has radio buttons for 'Port' (selected) and 'Trunk'. Below this, there are input fields for 'Protocol Group ID' (set to 1), 'VLAN ID (1-4094)', and 'Priority (1-7)'. At the bottom right are 'Apply' and 'Revert' buttons.

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 100: Showing the Interface to Protocol Group Mapping

The screenshot shows the 'VLAN > Protocol' configuration page with the 'Action' set to 'Show'. The 'Interface' section has radio buttons for 'Port' (selected) and 'Trunk'. Below this, a table titled 'Port to Protocol Group Mapping Table: Total: 1' is displayed. The table has four columns: a checkbox, 'Protocol Group ID', 'VLAN ID', and 'Priority'. One row is shown with values 1, 2, and 8 respectively. At the bottom right are 'Delete' and 'Revert' buttons.

	Protocol Group ID	VLAN ID	Priority
☐	1	2	8

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.
- ◆ Packet forwarding within the IP subnet requires that the following conditions are satisfied:
 - Inbound packet must be an IP packet (EtherType 0x0800).
 - The interface must be configured as a member of the subnet VLAN.
 - The IP address and subnet mask of the inbound packet must fall within specified subnet VLAN

The above conditions mean that any ARP requests will not be forwarded within the subnet VLAN because they are EtherType 0x0806 and do not meet the requirements for an IP subnet VLAN.

Also, note that TCP/UDP port numbers are not considered when checking whether or not packets are forwarded within the subnet VLAN

- ◆ 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 101: Configuring IP Subnet VLANs

VLAN > IP Subnet

Action: Add ▾

IP Address: 192.168.1.0

Subnet Mask: 255.255.255.0

VLAN (1-4094): 10

Priority (0-7): 1

Apply Revert

To show the configured IP subnet VLANs:

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

Figure 102: Showing IP Subnet VLANs

VLAN > IP Subnet

Action: Show ▾

IP Subnet to VLAN Mapping Table Total: 1

	IP Address	Subnet Mask	VLAN	Priority
☐	192.168.1.0	255.255.255.0	10	0

Delete Revert

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 103: Configuring MAC-Based VLANs

VLAN > MAC-Based

Action: Add ▼

MAC Address: 00-ab-cd-11-22-33 (00-XX-XX-XX-XX-XX OF XXXXXXXXXX)

Mask: (00-XX-XX-XX-XX-XX OF XXXXXXXXXX)

VLAN (1-4094): 10

Priority (0-7):

Apply Reset

To show the MAC addresses mapped to a VLAN:

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

Figure 104: Showing MAC-Based VLANs

VLAN > MAC-Based

Action: Show ▼

MAC-Based VLAN List Total: 1

	MAC Address	Mask	VLAN	Priority
1	00-AB-CD-11-22-33	FF-FF-FF-FF-FF-FF	10	0

Delete Reset

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 207](#)), the target port cannot be set to the same target ports as that used for port mirroring (see [“Configuring Local Port Mirroring” on page 142](#)).
- ◆ 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: 1-28/54)

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 105: Configuring VLAN Mirroring

The screenshot shows a web interface for configuring VLAN mirroring. The title bar reads "VLAN > Mirror". Below the title bar, there is an "Action:" label followed by a dropdown menu showing "Add". Below this, there are two rows of configuration fields. The first row is labeled "Source VLAN" and has a dropdown menu showing "1". The second row is labeled "Target Port" and has a dropdown menu showing "2". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

To show the VLANs to be mirrored:

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

Figure 106: Showing the VLANs to Mirror



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 107: 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.
- ◆ **Direction** - Specifies to apply VLAN translation to ingress or egress traffic, or both.
- ◆ **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 108: Configuring VLAN Translation

VLAN > Translation

Action: Add

Interface: ☒ Port ☐ Trunk

Port: 2

Direction: Both

Old VLAN (1-4094):

New VLAN (1-4094):

Apply Reset

To show the mapping entries for VLANs translation:

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

Figure 109: Showing the Entries for VLAN Translation



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:

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

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 110: Displaying the Dynamic MAC Address Table



The screenshot shows the 'MAC Address > Dynamic' web interface. At the top, there is a tab labeled 'Show Dynamic MAC'. Below this, the 'Query by:' section includes a 'Sort Key' dropdown set to 'MAC Address', a search input field containing '00-00-00-00-00-00', and radio buttons for 'MAC Address', 'VLAN', and 'Interface'. The 'Interface' section has 'Port' and 'Trunk' dropdowns. A 'Query' button is located below the search fields. The main area displays a table titled 'Dynamic MAC Address List' with 7 total entries. The table has columns for MAC Address, VLAN, Interface, Type, and Life Time.

MAC Address	VLAN	Interface	Type	Life Time
00-60-E1-81-84-4A	1	Unit 1 / Port 1	Learn	Delete on Timeout
00-E0-0C-10-90-91	1	Unit 1 / Port 1	Learn	Delete on Timeout
00-E0-4C-66-12-66	1	Unit 1 / Port 1	Learn	Delete on Timeout
00-E0-4C-66-14-79	1	Unit 1 / Port 1	Learn	Delete on Timeout
CC-37-AB-A1-06-C0	1	Unit 1 / Port 3	Learn	Delete on Timeout
CC-37-AB-A1-06-C3	1	Unit 1 / Port 3	Learn	Delete on Timeout
00-37-A0-2B-DE-50	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 111: Clearing Entries in the Dynamic MAC Address Table



The screenshot shows a web interface for configuring MAC Address settings. The breadcrumb trail at the top is "MAC Address > Dynamic". Below this, there is a "Step:" label followed by a dropdown menu currently set to "Clear Dynamic MAC". Underneath, there is a "Clear by:" label followed by a dropdown menu set to "All". A "Clear" button is located at the bottom right of the form area.

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: 5-514 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 112: Setting the Address Aging Time



The screenshot shows a web interface for configuring MAC address settings. The title bar reads "MAC Address > Dynamic". Below the title bar, there is a dropdown menu labeled "Action:" with "Configure Aging" selected. The main configuration area has two sections: "Aging Status" with a checked checkbox labeled "Enabled", and "Aging Time (1-514)" with a text input field containing "300" and a unit label "sec". At the bottom right of the configuration area are two buttons: "Apply" and "Revert".

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 205](#)) 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 352](#)).
 - Security Status (see [“Configuring Port Security” on page 348](#)) is enabled on the same interface.

Parameters

These parameters are displayed:

- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Port** – Port Identifier. (Range: 1-28/54)

- ◆ **Trunk** – Trunk Identifier. (Range: 1-28)
- ◆ **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 113: Configuring MAC Address Learning

Port	Status
1	✓ Enabled
2	✓ Enabled
3	✓ Enabled
4	✓ Enabled
5	✓ Enabled
6	✓ Enabled
7	✓ 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 114: Configuring Static MAC Addresses

The screenshot shows a web interface for configuring static MAC addresses. The title bar reads "MAC Address > Static". Below the title bar, there is a section for "Action" with a dropdown menu set to "Add". The main configuration area contains the following fields:


- VLAN**: A dropdown menu with the value "1" selected.
- Interface**: A section with two radio buttons, "Port" (selected) and "Trunk".
- MAC Address**: A text input field containing "00-12-cf-94-34-da". To the right of the input field is a placeholder text: "(xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx)".
- Static Status**: A dropdown menu with the value "Permanent" selected.

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

To show the static addresses in MAC address table:

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

Figure 115: Displaying Static MAC Addresses



The screenshot shows a web interface for configuring static MAC addresses. At the top, there's a breadcrumb 'MAC Address > Static' and an 'Actions' menu with a 'Show' button. Below this is a table titled 'Static MAC Address to Interface Mapping Table' with a 'Total 2' indicator. The table has six columns: MAC Address, VLAN, Interface, Type, and Life Time. There are two rows of data. The first row shows MAC address '99-88-00-90-88-F0' in VLAN 1, mapped to the CPU interface, with a 'CPU' type and a 'Delete on Reboot' life time. The second row shows MAC address '00:12:CF:84:3A:D6' in VLAN 1, mapped to 'Unit 1 / Port 1', with a 'Config' type and a 'Permanent' life time. At the bottom of the table are 'Delete' and 'Revert' buttons.

MAC Address	VLAN	Interface	Type	Life Time
99-88-00-90-88-F0	1	CPU	CPU	Delete on Reboot
00:12:CF:84:3A:D6	1	Unit 1 / Port 1	Config	Permanent

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 211](#)).
- ◆ When mirroring VLAN traffic (see [“Configuring VLAN Mirroring” on page 195](#)) 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 142](#)).
- ◆ 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: 1-28/54)

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 116: Mirroring Packets Based on the Source MAC Address

MAC Address > Mirror

Action: Add

Source MAC: (00-XX-XX-XX-XX-XX or XXXXXXXXXXXX)

Target Port: 1

Apply Revert

To show the MAC addresses to be mirrored:

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

Figure 117: 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	1.0

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: 1-28/54)
- ◆ **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 118: Issuing MAC Address Traps (Global Configuration)

The screenshot shows the 'MAC Address > MAC Notification' web interface. At the top, there is a breadcrumb trail 'MAC Address > MAC Notification'. Below it, a 'Step:' dropdown menu is set to '1. Configure Global'. The main configuration area contains two settings: 'MAC Notification Traps' with a checkbox labeled 'Enabled' (which is currently unchecked), and 'MAC Notification Trap Interval (1-3600)' with a text input field containing the value '1' and a 'sec' label. At the bottom right of the form, there are two buttons: 'Apply' and 'Revert'.

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 119: Issuing MAC Address Traps (Interface Configuration)

MAC Address > MAC Notification

Step: 2. Configure Interface

Interface: ☒ Port ☐ Trunk

Port List: Total: 20

Port	MAC Notification Trap
1	☐ Enabled
2	☐ Enabled
3	☐ Enabled
4	☐ 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 120: STP Root Ports and Designated Ports




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

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


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

Figure 121: 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 229](#)). 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 122: 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 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: Disabled)
- ◆ **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 shut down the interface. (Default: Shutdown)
- ◆ **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 123: Configuring Port Loopback Detection

Spanning Tree > Loopback Detection

Interface: ☒ Port ☐ Trunk

Loopback Detection Port List: Total 20

Port	Status	Trap	Release Mode	Release	Action	Shutdown Interval (80-85400 sec)
1	☐ Enabled	☐ Enabled	Auto	Release	Shutdown	60
2	☐ Enabled	☐ Enabled	Auto	Release	Shutdown	60
3	☐ Enabled	☐ Enabled	Auto	Release	Shutdown	60
4	☐ Enabled	☐ Enabled	Auto	Release	Shutdown	60
5	☐ Enabled	☐ Enabled	Auto	Release	Shutdown	60
6	☐ Enabled	☐ Enabled	Auto	Release	Shutdown	60
7	☐ Enabled	☐ Enabled	Auto	Release	Shutdown	60
8	☐ Enabled	☐ Enabled	Auto	Release	Shutdown	60
9	☐ Enabled	☐ Enabled	Auto	Release	Shutdown	60
10	☐ Enabled	☐ Enabled	Auto	Release	Shutdown	60

Apply Revert

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 Protocol⁵

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 Protocol⁵

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.

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

- RSTP Mode – If RSTP is using 802.1D BPDUs on a port and receives an RSTP BPDU after the migration delay expires, RSTP restarts the migration delay timer and begins using RSTP BPDUs on that port.

◆ **Multiple Spanning Tree Protocol**

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

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

Parameters

These parameters are displayed:

Basic Configuration of Global Settings

- ◆ **Spanning Tree Status** – Enables/disables STA on this switch.
(Default: Disabled)
- ◆ **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 how the system floods BPDUs to other ports when spanning tree is disabled globally on the switch or disabled on specific ports.
 - To VLAN: Floods BPDUs to all other spanning tree-disabled 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 spanning tree-disabled 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 number of BPDU transmissions permitted within the Hello Time interval. (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**⁶ – The revision for this MSTI. (Range: 0-65535; Default: 0)
- ◆ **Region Name**⁶ – 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.

6. The MST name and revision number 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.
5. Click Apply

Figure 124: Configuring Global Settings for STA (STP)

Spanning Tree > STA

Steps: 1. Configure Global Actions: Configure

Spanning Tree Status ☒ Enabled

Spanning Tree Type STP

Priority (0-61440, in steps of 4096)

BPDU Flooding To VLAN

Cisco Prestandard Status ☐ Enabled

Advanced:

Path Cost Method Long

Transmission Limit (1-10)

When the Switch Becomes Root:

Hello Time (1-10) sec

Maximum Age (5-40) sec

Forward Delay (4-30) sec

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

Apply Revert

Figure 125: 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 126: Configuring Global Settings for STA (MSTP)

Spanning Tree > STA

Step: 1. Configure Global Action: Configure

Spanning Tree Status ☒ Enabled

Spanning Tree Type MSTP

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

MSTP Configuration

Max Instance Numbers 64

Configuration Digest 0xAC36177F50283CD4B63621D64826DE2

Region Revision (0-65535) 0

Region Name 00:00:00:00:00:FD

Max Hop Count (1-40) 20

Apply Revert

Displaying Global Settings for STA

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

Parameters

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


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



The screenshot shows the 'Spanning Tree > STA' configuration page. At the top, there are tabs for 'Steps' (1. Configure Global) and 'Actions' (Show Information). Below this is a section titled 'Spanning Tree Information' containing two columns of settings.

Spanning Tree Information	
Spanning Tree Status	Enabled
Designated Root	32768.000000000000000000
Root Port	0
Root Path Cost	0
Configuration Changes	2
Last Topology Change	0 days, 3 hours, 51 minutes, 11 seconds
Spanning Tree Type	RSTP
Bridge ID	32768.000000000000000000
Max Age	20 sec
Hello Time	2 sec
Forward Delay	15 sec

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 or when spanning tree is disabled on a specific port. When flooding is enabled, BPDUs are flooded to all other spanning-tree disabled ports on the switch or to all other spanning-tree disabled ports within the receiving port’s native VLAN, as specified by the Spanning Tree BPDU Flooding attribute ([page 217](#)). (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 method⁷, 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

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 128: Determining the Root Port



For BPDUs received by i1 on SW3, the path cost is 0.
For BPDUs received by i2 on SW3, the path cost is that of i1 on SW2.

7. Refer to “Configuring Global Settings for STA” on page 215 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).

The root path cost for i1 on SW3 used to compete for the role of root port is $0 + \text{path cost of i1 on SW3}$; 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 $10000 + \text{path cost of i2 on SW3}$.

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

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

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

- ◆ **Admin Link Type** – The link type attached to this interface.
 - Point-to-Point – A connection to exactly one other bridge.
 - Shared – A connection to two or more bridges.
 - Auto – The switch automatically determines if the interface is attached to a point-to-point link or to shared media. (This is the default setting.)
- ◆ **Root Guard** – STA allows a bridge with a lower bridge identifier (or same identifier and lower MAC address) to take over as the root bridge at any time. Root Guard can be used to ensure that the root bridge is not formed at a suboptimal location. Root Guard should be enabled on any designated port connected to low-speed bridges which could potentially overload a slower link by taking over as the root port and forming a new spanning tree topology. It could also be used to form a border around part of the network where the root bridge is allowed. (Default: Disabled)
- ◆ **Admin Edge Port** – Since end nodes **cannot** cause forwarding loops, they can pass directly through to the spanning tree forwarding state. Specifying Edge Ports provides quicker convergence for devices such as workstations or servers, retains the current forwarding database to reduce the amount of frame flooding required to rebuild address tables during reconfiguration events, does not cause the spanning tree to initiate reconfiguration when the interface changes state, and also overcomes other STA-related timeout problems. However, remember that Edge Port should only be enabled for ports connected to an end-node device. (Default: Auto)
 - **Enabled** – Manually configures a port as an Edge Port.
 - **Disabled** – Disables the Edge Port setting.
 - **Auto** – The port will be automatically configured as an edge port if the edge delay time expires without receiving any RSTP or MSTP BPDUs. Note that edge delay time (802.1D-2004 17.20.4) equals the protocol migration time if a port's link type is point-to-point (which is 3 seconds as defined in

IEEE 802.3D-2004 17.20.4); otherwise it equals the spanning tree's maximum age for configuration messages (see maximum age under ["Configuring Global Settings for STA" on page 215](#)).

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

- If spanning tree mode is set to STP ([page 215](#)), edge-port mode cannot automatically transition to operational edge-port state using the automatic setting.
- If loopback detection is enabled ([page 213](#)) 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 226](#)).

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 Guard Auto Recovery** – Automatically re-enables an interface after the specified interval. (Default: Disabled)
- ◆ **BPDU Guard Auto Recovery Interval** – The time to wait before re-enabling an interface. (Range: 30-86400 seconds; Default: 300 seconds)
- ◆ **BPDU Filter** – BPDU filtering allows you to avoid transmitting BPDUs on configured edge ports that are connected to end nodes. By default, STA sends BPDUs to all ports regardless of whether administrative edge is enabled on a port. BPDU filtering is configured on a per-port basis. (Default: Disabled)

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

- ◆ **Migration** – If at any time the switch detects STP BPDUs, including Configuration or Topology Change Notification BPDUs, it will automatically set

◆ **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 222](#).
- ◆ **Oper Edge Port** – This parameter is initialized to the setting for Admin Edge Port in STA Port Configuration on [page 222](#) (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 130: 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 131: Displaying Interface Settings for STA

The screenshot shows the 'Spanning Tree > STA' configuration page. At the top, there's a 'Step' dropdown set to '2. Configure Interface' and an 'Action' dropdown set to 'Show Information'. Below this, there are radio buttons for 'Port' (selected) and 'Trunk'. A 'Spanning Tree Port List' section shows a table with 11 columns: 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 table lists four ports (1, 2, 3, 4) with their respective settings.

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	1	0	32768.E490731B5E00	128.1	10000	Point-to-Point	Disabled	Disabled
2	Enabled	Enabled	Discarding	0	0	32768.E490731B5E00	128.2	10000	Point-to-Point	Disabled	Disabled
3	Enabled	Enabled	Discarding	0	0	32768.E490731B5E00	128.3	10000	Point-to-Point	Disabled	Disabled
4	Enabled	Enabled	Discarding	0	0	32768.E490731B5E00	128.4	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 215) 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 215).
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 132: Creating an MST Instance

The screenshot shows a web interface titled "Spanning Tree > MSTP". Below the title, there are two dropdown menus: "Steps" set to "1. Configure Global" and "Actions" set to "Add". Below these, there are three input fields: "MST ID (0-4094)" with the value "1", "VLAN ID (1-4094)" with the value "1", and "Priority (0-61440, in steps of 4096)" which is empty. At the bottom right, there are two buttons: "Apply" and "Revert".

To show the MSTP instances:

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


Figure 133: Displaying MST Instances



To modify the priority for an MST instance:

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

Figure 134: Modifying the Priority for an MST Instance



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 221.

Figure 135: Displaying Global Settings for an MST Instance

The screenshot shows the 'Spanning Tree > MSTP' configuration page. At the top, there is a 'Step:' dropdown menu set to '1. Configure Global' and an 'Action:' dropdown menu set to 'Show information'. Below this, the 'MST ID' is set to '1'. The page displays a table of global settings for this instance:

Priority	0	Designated Root	32768:0030F124568D
Bridge ID	20	Root Port	2
Max Age	15 sec	Root Path Cost	32768:0000161001D
Hello Time	20 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 136: Adding a VLAN to an MST Instance

The screenshot shows the 'Spanning Tree > MSTP' configuration page with the 'Action:' dropdown menu set to 'Add Member'. The 'MST ID' is set to '1'. Below it, the 'VLAN ID (1-4094)' field 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 137: Displaying Members of an MST Instance



Configuring Interface Settings for MSTP

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

Parameters

These parameters are displayed:

- ◆ **MST ID** – Instance identifier to configure. (Default: 0)
- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **STA Status** – Displays the current state of this interface within the Spanning Tree. (See [“Displaying Interface Settings for STA” on page 226](#) 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 215), 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 223.

The default path costs are listed in Table 12 on page 223.

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 138: Configuring MSTP Interface Settings

Port	STA Status	Priority (0-240, in steps of 16)	Admin MST Path Cost (0-268435456, 0: Auto)
1	Forwarding	128	0
2	Discarding	128	0
3	Discarding	128	0
4	Discarding	128	0

To display MSTP parameters for a port or trunk:

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

Figure 139: Displaying MSTP Interface Settings

Spanning Tree > MSTP

Step: 2. Configure Interface Action: Show Information

MST ID: 1

Interface: ☒ Port ☐ Trunk

Spanning Tree Port List Total: 20

Port	STA Status	Forward Transitions	Designated Cost	Designated Bridge	Designated Port	Oper Path Cost	Oper Link Type	Oper Edge Port	Port Role
1	Forwarding	1	8	32768.E480731B5E00	128.1	10000	Point-to-Point	Enabled	Disabled
2	Discarding	0	8	32768.E480731B5E00	128.2	10000	Point-to-Point	Enabled	Disabled
3	Discarding	0	8	32768.E480731B5E00	128.3	10000	Point-to-Point	Enabled	Disabled
4	Discarding	0	8	32768.E480731B5E00	128.4	10000	Point-to-Point	Enabled	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+)
- ◆ **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)

- ◆ **Resolution** – Indicates the resolution at which the rate can be configured.

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 140: Configuring Rate Limits

The screenshot shows the 'Traffic > Rate Limit' web interface. At the top, there are radio buttons for 'Interface' type: 'Port' (selected) and 'Trunk'. Below this is a 'Port Rate Limit List' with a 'Total: 28' indicator. The main part of the interface is a table with columns for 'Port', 'Type', 'Input' (Status and Rate), and 'Output' (Status and Rate). The table lists five ports, all of type '10GBASE-T', with their respective status and rate limit settings.

Port	Type	Input		Output	
		Status	Rate (kb/s/sec) (64-10000000)	Status	Rate (kb/s/sec) (64-10000000)
1	10GBASE-T	☐ Enabled	1000000	☐ Enabled	1000000
2	10GBASE-T	☐ Enabled	1000000	☐ Enabled	1000000
3	10GBASE-T	☐ Enabled	1000000	☐ Enabled	1000000
4	10GBASE-T	☐ Enabled	1000000	☐ Enabled	1000000
5	10GBASE-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.

- ◆ 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+)
- ◆ **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: disabled for broadcast, 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)
- ◆ **Resolution** – Indicates the resolution at which the rate can be configured.

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 141: Configuring Storm Control

Traffic > Storm Control

Interface: ☒ Port ☐ Trunk

Port Storm Control List Test 28

Port	Type	Unknown Unicast		Multicast		Broadcast	
		Status	Rate (packets/sec) (500-14251000)	Status	Rate (packets/sec) (500-14251000)	Status	Rate (packets/sec) (500-14251000)
1	100BASE-T	☒ Enabled	500	☒ Enabled	500	☒ Enabled	500
2	100BASE-T	☒ Enabled	500	☒ Enabled	500	☒ Enabled	500
3	100BASE-T	☒ Enabled	500	☒ Enabled	500	☒ Enabled	500
4	100BASE-T	☒ Enabled	500	☒ Enabled	500	☒ Enabled	500
5	100BASE-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 142: 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 243](#)).
- ◆ 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 143: 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 238](#)). 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 243](#)).

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: 1-300 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: 1-900 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: 1-300 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: 1-900 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 144: Configuring ATC Timers

The screenshot shows the 'Traffic > Auto Traffic Control' configuration page. At the top, there is a 'Step: 1. Configure Global' dropdown. Below this, there are four rows of configuration fields:

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

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

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 142 on page 240](#). 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 145: Configuring ATC Interface Attributes

Traffic > Auto Traffic Control

Step: 2. Configure Interface

Storm Control ☒ Broadcast ☐ Multicast

Auto Traffic Control Broadcast List Total: 20

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 v	☐ Enabled	128	128	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	Release
2	☐ Enabled	Rate Control v	☐ Enabled	128	128	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	Release
3	☐ Enabled	Rate Control v	☐ Enabled	128	128	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	Release
4	☐ Enabled	Rate Control v	☐ 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 146: Setting the Default Port Priority

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

Selecting the Queue Mode

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

Command Usage

- ◆ Strict priority requires all traffic in a higher priority queue to be processed before lower priority queues are serviced.
- ◆ WRR queuing specifies a relative weight for each queue. WRR uses a predefined relative weight for each queue that determines the percentage of service time the switch services each queue before moving on to the next queue. This prevents the head-of-line blocking that can occur with strict priority queuing.

- ◆ If Strict and WRR mode is selected, a combination of strict 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: 1-127; 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 queuing is selected, the queues which are serviced first must be specified by enabling strict mode parameter in the table.
5. Click Apply.

Figure 147: Setting the Queue Mode (Strict)

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

Figure 148: Setting the Queue Mode (WRR)

The screenshot shows the 'Traffic > Priority > Queue' configuration page. The 'Queue Mode' is set to 'WRR' in a dropdown menu. 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 'Queue ID' column lists values from 0 to 7. The 'Weight' column has input fields with values 1, 2, 4, 6, 8, 10, 12, and 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 149: Setting the Queue Mode (Strict and WRR)

The screenshot shows the 'Traffic > Priority > Queue' configuration page. The 'Queue Mode' is set to 'Strict and WRR' in a dropdown menu. Below the dropdown is a 'Queue Setting Table' with a 'Total: 8' label. The table has three columns: 'Queue ID', 'Strict Mode', and 'Weight (1-255)'. The 'Queue ID' column lists values from 0 to 7. The 'Strict Mode' column has dropdown menus, all set to 'Disabled'. The 'Weight' column has input fields with values 1, 2, 4, 6, 8, 10, 12, and 14. Below the table are 'Apply' and 'Revert' buttons.

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

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 257](#)).

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.

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 150: 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 an interface.

Figure 151: Showing CoS Values to Egress Queue Mapping

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

Layer 3/4 Priority Settings

Mapping Layer 3/4 Priorities to CoS Values

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

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



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

Setting Priority Processing to IP Precedence/DSCP or CoS

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

Command Usage

- ◆ If the QoS mapping mode is set to IP Precedence, and the ingress packet type is IPv4, then priority processing will be based on the IP Precedence value in the ingress packet.
- ◆ If the QoS mapping mode is set to DSCP, and the ingress packet type is IPv4, then priority processing will be based on the DSCP value in the ingress packet.
- ◆ If the QoS mapping mode is set to either IP Precedence or DSCP, and a non-IP packet is received, the packet's CoS and CFI (Canonical Format Indicator) values are used for priority processing if the packet is tagged. For an untagged packet, the default port priority (see [page 247](#)) 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 247](#)) is used for priority processing.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28/54)
- ◆ **Trust Mode**
 - **IP Precedence** – Maps layer 3/4 priorities using IP Precedence values.
 - **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 152: 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, 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, ingress-dscp = ingress-dscp10 * 10 + 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 153: Configuring DSCP to DSCP Internal Mapping

Traffic > Priority > DSCP to DSCP

Action: Configure

Port: 1

DSCP (0-63):

PHB (0-7):

Drop Precedence: 0: Green

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 154: Showing DSCP to DSCP Internal Mapping

Figure 154 shows the 'DSCP to DSCP Mapping List' table. The table has three columns: DSCP, PHB, and Drop Precedence. The data is as follows:

DSCP	PHB	Drop Precedence
0	0	0: Green
1	0	1: Red
2	0	0: Green
3	0	2: Yellow
4	0	0: Green
5	0	1: Red
6	0	0: Green
7	0	2: Yellow
8	1	0: Green
9	1	1: Red

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 258](#).
- ◆ 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 155: Configuring CoS to DSCP Internal Mapping

Traffic > Priority > CoS to DSCP

Action: Configure

Port: 1

CoS (0-7): 0

CFI (0-1): 1

PHB (0-7): 0

Drop Precedence: 0 Green

Apply 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 156: Showing CoS to DSCP Internal Mapping

Traffic > Priority > CoS to DSCP

Action: Show

Port: 1

CoS to DSCP Mapping List: Total: 16

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

Default Revert

Mapping IP Precedence Values to Internal DSCP Values

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

The Type of Service (ToS) octet in the IPv4 header includes three precedence bits defining eight different priority levels ranging from highest priority for network control packets to lowest priority for routine traffic. The default IP Precedence values map one-to-one to the Class of Service values (that is, Precedence value 0

maps to PHB value 0, and so forth). Bits 6 and 7 are used for network control, and the other bits for various application types. The ToS bits are defined in [Table 18](#).

Table 18: Mapping IP Precedence

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

Command Usage

- ◆ Enter per-hop behavior and drop precedence for any of the IP Precedence values 0 - 7.
- ◆ If the priority mapping mode is set the IP Precedence and the ingress packet type is IPv4, then the IP Precedence-to-PHB/Drop Precedence mapping table is used to generate priority and drop precedence values for internal processing.
- ◆ Random Early Detection starts dropping yellow and red packets when the buffer fills up to 0x60 packets, and then starts dropping any packets regardless of color when the buffer fills up to 0x80 packets.

Parameters

These parameters are displayed in the web interface:

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

Table 19: Default Mapping of IP Precedence to Internal PHB/Drop Values

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

Web Interface

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

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

Figure 157: Configuring IP Precedence to DSCP Internal Mapping

Traffic > Priority > IP Precedence to DSCP

Action: Configure

Port: 1

IP Precedence (0-7): 7

PHB (0-7): 5

Drop Precedence: 0: Green

Apply Revert

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

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

Figure 158: Showing the IP Precedence to DSCP Internal Map

Traffic > Priority > IP Precedence to DSCP

Action: Show

Port: 1

IP Precedence to DSCP Mapping List Total: 8

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

Clear

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 264](#)).

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 16 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 to create a service policy for a specific interface that defines packet classification, service tagging, and bandwidth policing. Note that one or more class maps can be assigned to a policy map.
- ◆ Up to 32 class maps can be configured.

Parameters

These parameters are displayed:

Add

- ◆ **Class Name** – Name of the class map. (Range: 1-32 characters)
- ◆ **Type** – Only one match command is permitted per class map, so the match-any field refers to the lone criteria specified by this command.

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

Add Rule

- ◆ **Class Name** – Name of the class map.
- ◆ **Type** – Only one match command is permitted per class map, so the match-any field refers to the lone criteria specified by this command.
- ◆ **ACL** – Name of an access control list. Any type of ACL can be specified, including standard or extended IPv4/IPv6 ACLs and MAC ACLs.
- ◆ **IP DSCP** – A DSCP value. (Range: 0-63)
- ◆ **IP Precedence** – An IP Precedence value. (Range: 0-7)
- ◆ **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: 1-28/54)

Web Interface

To configure a class map:

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


Figure 159: Configuring a Class Map

The screenshot shows a web interface for configuring a class map. At the top, it says 'Traffic > DiffServ'. Below that, there's a 'Step:' dropdown menu set to '1. Configure Class' and an 'Action:' dropdown menu set to 'Add'. The main form has three fields: 'Class Name' with the value 'rs-class', 'Type' with a dropdown menu set to '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 160: Showing Class Maps



To edit the rules for a class map:

1. Click Traffic, DiffServ.
2. Select Configure Class from the Step list.
3. Select Add Rule from the Action list.
4. Select the name of a class map.
5. Specify type of traffic for this class based on an access list, a DSCP or IP Precedence value, 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 161: Adding Rules to a Class Map



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 162: Showing the Rules for a Class Map



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 264](#)), modify service tagging, and enforce bandwidth policing. A policy map can then be bound by a service policy to one or more interfaces ([page 277](#)).

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 CoS value, or a source port. A policy map must then be configured to indicate 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 $T_p(0) = BP$ and the token count $T_c(0) = BC$. Thereafter, the token count T_p 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 277](#)). 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 277](#)) 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.
- ◆ **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 258](#)).

- **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 256](#)).

- **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) – 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) – 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 rate (CIR) 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)** – 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)** – 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 maximum rate (CIR) 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 maximum rate (CIR) but is within the excess burst size (BE) will be dropped or the DSCP service level will be reduced.

- **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63)

- **Drop** – Drops out of conformance traffic.

- **Violate** – Specifies whether the traffic that exceeds the excess burst size (BE) will be dropped or the DSCP service level will be reduced.

- **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63)
- **Drop** – Drops out of conformance traffic.
- **trTCM (Police Meter)** – Defines the committed information rate (CIR, or maximum throughput), peak information rate (PIR), and their associated burst sizes – committed burst size (BC, or burst rate) and peak burst size (BP), and the action to take for traffic conforming to the maximum throughput, exceeding the maximum throughput but within the peak information rate, or exceeding the peak information rate. In addition to the actions defined by this command to transmit, remark the DSCP service value, or drop a packet, the switch will also mark the two color bits used to set the drop precedence of a packet.

The color modes include “Color-Blind” which assumes that the packet stream is uncolored, and “Color-Aware” which assumes that the incoming packets are pre-colored. The functional differences between these modes is described at the beginning of this section under “trTCM Police Meter.”

- **Committed Information Rate (CIR)** – 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)** – 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: 0-10000000 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 maximum rate (CIR) 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 maximum rate (CIR) but is within the peak information rate (PIR) will be dropped or the DSCP service level will be reduced.

- **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63).
- **Drop** – Drops out of conformance traffic.
- **Violate** – Specifies whether the traffic that exceeds the peak information rate (PIR) will be dropped or the DSCP service level will be reduced.
- **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63).
- **Drop** – Drops out of conformance traffic.

Web Interface

To configure a policy map:

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

Figure 163: Configuring a Policy Map

The screenshot shows a web interface for configuring a policy map. At the top, it says "Traffic > DiffServ". Below that, there are two dropdown menus: "Step:" with "2. Configure Policy" selected, and "Action:" with "Add" selected. Below these are two text input fields: "Policy Name" with the value "fd-policy" and "Description" with the value "for the software group". At the bottom right, there are two buttons: "Apply" and "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 164: Showing Policy Maps



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 165: Adding Rules to a Policy Map

Traffic > DiffServ

Steps: 2. Configure Policy Action: Add Rule

Policy Name: nt-policy

Rule:

Class Name: nt-class

☒ Action: Set P-B (D-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

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 166: Showing the Rules for a Policy Map

Traffic > DiffServ

Steps: 2. Configure Policy Action: Show Rule

Policy Name: nt-policy

Rule List: Total: 1

	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	Violate
☐	nt-class	Set P-B 3	Flow	1000000	4000				Transmit		Drop

Delete Save

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 167: Attaching a Policy Map to a Port

The screenshot shows the 'Traffic > DiffServ' web interface. At the top, there is a 'Step' dropdown menu set to '3. Configure Interface'. Below this is a table titled 'Port Service Policy List' with columns for 'Port', 'Ingress', and 'Egress'. The table has 10 rows, numbered 1 to 10. Each row has a 'Port' column with a dropdown menu, and 'Ingress' and 'Egress' columns with checkboxes. At the bottom of the table, there are 'Apply' and 'Revert' buttons.

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

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 167](#)).

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 168: Configuring a Voice VLAN

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 169: Configuring an OUI Telephony List

Traffic > VoIP

Step: 2. Configure OUI Actions: Add

Telephony OUI: 00-e0-bb-00-00-00 (XX-XX-XX-XX-XX-XX of XXXXXXXXXXXXX)

Mask: 8-8-8-00-00-00 (XX-XX-XX-XX-XX-XX of XXXXXXXXXXXXX)

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 170: Showing an OUI Telephony List

Traffic > VoIP

Step: 2. Configure OUI Actions: Show

Telephony OUI List - Total: 3

	Telephony OUI	Mask	Description
☐	88-88-88-00-00-00	FF-FF-FF-00-00-00	old phones
☐	00-11-22-33-44-55	FF-FF-FF-00-00-00	new phones
☐	00-88-88-00-00-00	FF-FF-FF-FF-FF-FF	old phones

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 167](#)).

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 416](#) 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 171: Configuring Port Settings for a Voice VLAN



The screenshot shows the 'Traffic > VoIP' configuration page. The 'Steps' dropdown is set to 'Configure Interface'. Below the table, there are icons for help, refresh, and delete. The table has 6 columns: Port, Mode, Security, Discovery Protocol, Priority (1-4), and Remaining Age (minutes).

Port	Mode	Security	Discovery Protocol	Priority (1-4)	Remaining Age (minutes)
1	None	☐ Enabled	☒ OUI ☐ LLDP	5	NA
2	None	☐ Enabled	☒ OUI ☐ LLDP	5	NA
3	None	☐ Enabled	☒ OUI ☐ LLDP	5	NA
4	None	☐ Enabled	☒ OUI ☐ LLDP	5	NA
5	None	☐ Enabled	☒ OUI ☐ LLDP	5	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).
- ◆ **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.
- ◆ **DHCP Snooping** – Filter IP traffic on insecure ports for which the source address cannot be identified via DHCP snooping.
- ◆ **DHCPv6 Snooping** – Filter IPv6 traffic on insecure ports for which the source address cannot be identified via DHCPv6 snooping.

- ◆ **ND Snooping** – Maintains an IPv6 prefix table and user address binding table. These tables can be used for stateless address auto-configuration or for address filtering by IPv6 Source Guard.
- ◆ **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.
- ◆ **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.
- ◆ **Application Filter** – Discards CDP or PVST packets.



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 287](#).
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 172: Configuring the Authentication Sequence




The screenshot shows a web interface with a breadcrumb trail at the top: "Security > AAA > System Authentication". Below this, there is a label "Authentication Sequence" followed by a dropdown menu. The dropdown menu is currently set to "Local, RADIUS". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

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 173: 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: 1-10; Default: 5)
- **Authentication Retries** – Number of times the switch tries to authenticate logon access via the authentication server. (Range: 1-5; Default: 3)
- **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 one of five TACACS+ 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 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: 1-10; Default: 5)
- **Authentication Retries** – Number of times the switch tries to authenticate logon access via the authentication server. (Range: 1-5; Default: 3)
- **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)

When specifying the priority sequence for a sever, the server index must already be defined (see [“Configuring Local/Remote Logon Authentication” on page 287](#)).

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 174: Configuring Remote Authentication Server (RADIUS)

Security > AAA > Server

Step: 1. Configure Server

Server Type: ☒ RADIUS ☐ TACACS+

☐ Global | Server Index: ☒ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5

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: [REDACTED]

Confirm Authentication Key: [REDACTED]

Apply Revert

Figure 175: Configuring Remote Authentication Server (TACACS+)

Security > AAA > Server

Step: 1. Configure Server

Server Type: ☐ RADIUS ☒ TACACS+

☐ Global | Server Index: ☒ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5

Server IP Address: [REDACTED]

Authentication Server TCP Port (1-65535): [REDACTED]

Authentication Timeout (1-640): [REDACTED] sec

Authentication Retries (1-30): [REDACTED]

☐ Set Key

Authentication Key: [REDACTED]

Confirm Authentication Key: [REDACTED]

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 176: Configuring AAA Server Groups

Security > AAA > Server

Steps: 2. Configure Group Actions: Add

Server Type ☒ RADIUS ☐ TACACS+

RADIUS Group Name

radius

Sequence At Priority 1

1

Sequence At Priority 2

3

Sequence At Priority 3

5

Sequence At Priority 4

2

Sequence At Priority 5

None

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 177: Showing AAA Server Groups

Security > AAA > Server

Steps: 2. Configure Group Actions: Show

Server Type ☒ RADIUS ☐ TACACS+

RADIUS Group List Total: 3

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

Delete

Revert

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 287](#)). 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 178: Configuring Global Settings for AAA Accounting



The screenshot shows a web interface for configuring AAA Accounting. The breadcrumb navigation at the top reads "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 "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 179: Configuring AAA Accounting Methods

Security > AAA > Accounting

Step: 2. Configure Method Action: Add

Accounting Type: RADIUS

Method Name: default

Accounting Notice: Start-Stop

Server Group Name: radius

Apply 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 180: Showing AAA Accounting Methods

Security > AAA > Accounting

Step: 2. Configure Method Action: Show

Method List Total: 10

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

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 181: 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: 20

Port	Method Name
1	
2	
3	
4	
5	
6	
7	

Figure 182: Configuring AAA Accounting Service for Command Service

Security > AAA > Accounting

Step: 3. Configure Service

Accounting Type: ☐ 802.1X ☒ Command ☐ EXEC

Command Method List Total: 10

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

Figure 183: 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 184: Displaying a Summary of Applied AAA Accounting Methods

Accounting Type	Method Name	Server Group Name	Interface
RADIUS	default	default	
Command 0	default	default	
Command 1	default	default	
Command 2	default	default	
Command 3	default	default	
Command 4	default	default	
Command 5	default	default	
Command 6	default	default	
Command 7	default	default	
Command 8	default	default	

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 185: Displaying Statistics for AAA Accounting Sessions

User Name	Accounting Type	Interface	Time Elapsed
Bob	RADIUS	default	3:44:55
Ted	RADIUS	default	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:
 - **Command** – Administrative authorization to apply to commands entered at specific CLI privilege levels.
 - **Exec** – 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 287](#)). 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 186: 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 187: Showing AAA Authorization Methods

Security > AAA > Authorization

Step: 1. Configure Method Action: Show

Method List Type: 2

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

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 188: 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 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 189: 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 or Modify from the Action list.
3. Specify or select a user name, select the user's access level, then enter a password if required and confirm it.
4. Click Apply.

Figure 190: Configuring User Accounts

The screenshot shows a web interface titled "Security > User Accounts". At the top, there is a tab labeled "Action: Add". Below this, the form contains the following fields:

- User Name:** A text input field containing the value "bob".
- Access Level:** A dropdown menu with a downward arrow, currently showing the value "5".
- Password Type:** A dropdown menu with a downward arrow, currently showing the value "Encrypted Password".
- Password:** A text input field containing a masked password (represented by asterisks).
- Confirm Password:** A text input field containing a masked password (represented by asterisks).

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

To show user accounts:

1. Click Security, User Accounts.
2. Select Show from the Action list.

Figure 191: Showing User Accounts

Security > User Accounts		
Actions: Show		
User Account List Total: 3		
	User Name	Access Level
☐	admin	15
☐	guest	0
☐	sub	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 287.)

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 192: Configuring Global Settings for Web Authentication

The screenshot shows the 'Security > Web Authentication' configuration page. At the top, there is a breadcrumb trail 'Security > Web Authentication' and a 'Step:' dropdown menu set to '1. Configure Global'. Below this, the 'Web Authentication Status' is set to 'Enabled' with a checkbox. The 'Session Timeout (300-3600)' is set to '3600' seconds. The 'Quiet Period (1-180)' is set to '60' seconds. The 'Login Attempts (1-3)' is set to '3'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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.
- ◆ **Revert** – Restores the previous configuration settings.

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 193: Configuring Interface Settings for Web Authentication

Authenticated Host List Total: 2	
Host IP Address	Remaining Session Time (sec)
192.1.1.1	300
192.2.2	100

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 288.)

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]
 - **Radius:Avenda Avenda-Tag-ID** = 0

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 20: Dynamic QoS Profiles

Profile	Attribute Syntax	Example
DiffServ	service-policy-in = <i>policy-map-name</i>	service-policy-in=p1
Rate Limit	rate-limit-input = <i>rate</i>	rate-limit-input=100 (kbps)
	rate-limit-output = <i>rate</i>	rate-limit-output=200 (kbps)
802.1p	switchport-priority-default = <i>value</i>	switchport-priority-default=2
IP ACL	ip-access-group-in = <i>ip-acl-name</i>	ip-access-group-in=ipv4acl
IPv6 ACL	ipv6-access-group-in = <i>ipv6-acl-name</i>	ipv6-access-group-in=ipv6acl
MAC ACL	mac-access-group-in = <i>mac-acl-name</i>	mac-access-group-in=macAcl

- ◆ Multiple profiles can be specified in the Filter-ID attribute by using a semicolon to separate each profile.

For example, the attribute “service-policy-in=pp1;rate-limit-input=100” specifies that the diffserv profile name is “pp1,” and the ingress rate limit profile value is 100 kbps.

- ◆ If duplicate profiles are passed in the Filter-ID attribute, then only the first profile is used.

For example, if the attribute is “service-policy-in=p1;service-policy-in=p2,” then the switch applies only the DiffServ profile “p1.”

- ◆ Any unsupported profiles in the Filter-ID attribute are ignored.

For example, if the attribute is “map-ip-dscp=2:3;service-policy-in=p1,” then the switch ignores the “map-ip-dscp” profile.

- ◆ When authentication is successful, the dynamic QoS information may not be passed from the RADIUS server due to one of the following conditions (authentication result remains unchanged):

- The Filter-ID attribute cannot be found to carry the user profile.
- The Filter-ID attribute is empty.
- The Filter-ID attribute format for dynamic QoS assignment is unrecognizable (can not recognize the whole Filter-ID attribute).

- ◆ Dynamic QoS assignment fails and the authentication result changes from success to failure when the following conditions occur:

- Illegal characters found in a profile value (for example, a non-digital character in an 802.1p profile value).
- Failure to configure the received profiles on the authenticated port.

- ◆ When the last user logs off on a port with a dynamic QoS assignment, the switch restores the original QoS configuration for the port.

- ◆ When a user attempts to log into the network with a returned dynamic QoS profile that is different from users already logged on to the same port, the user is denied access.
- ◆ While a port has an assigned dynamic QoS profile, any manual QoS configuration changes only take effect after all users have logged off the port.

Configuring Global Settings for Network Access

MAC address authentication is configured on a per-port basis, however there are two configurable parameters that apply globally to all ports on the switch. Use the Security > Network Access (Configure Global) page to configure MAC address authentication aging and reauthentication time.

Parameters

These parameters are displayed:

- ◆ **Aging Status** – Enables aging for authenticated MAC addresses stored in the secure MAC address table. (Default: Disabled)

This parameter applies to authenticated MAC addresses configured by the MAC Address Authentication process described in this section, as well as to any secure MAC addresses authenticated by 802.1X, regardless of the 802.1X Operation Mode (Single-Host, Multi-Host, or MAC-Based authentication as described on [page 353](#)).

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 194: Configuring Global Settings for Network Access

Security > Network Access

Step: 1. Configure Global

Aging Status: ☒ Enabled

Reauthentication Time (120-1000000): 30000 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 Count**⁸ – Sets the maximum number of MAC addresses that can be authenticated on a port via MAC authentication; that is, the Network Access process described in this section. (Range: 1-1024; Default: 1024)

- ◆ **Network Access Max MAC Count**⁸ – 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 165](#)). 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 353](#)).

A port can only be assigned to the guest VLAN in case of failed authentication, and switchport mode is set to Hybrid. (See [“Adding Static Members to VLANs” on page 167](#).)

8. 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.

- ◆ **Dynamic VLAN** – Enables dynamic VLAN assignment for an authenticated port. When enabled, any VLAN identifiers returned by the RADIUS server through the 802.1X authentication process are applied to the port, providing the VLANs have already been created on the switch. (GVRP is not used to create the VLANs.) (Default: Enabled)

The VLAN settings specified by the first authenticated MAC address are implemented for a port. Other authenticated MAC addresses on the port must have the same VLAN configuration, or they are treated as authentication failures.

If dynamic VLAN assignment is enabled on a port and the RADIUS server returns no VLAN configuration (to the 802.1X authentication process), the authentication is still treated as a success, and the host is assigned to the default untagged VLAN.

When the dynamic VLAN assignment status is changed on a port, all authenticated addresses mapped to that port are cleared from the secure MAC address table.

- ◆ **Dynamic QoS** – Enables dynamic QoS assignment for an authenticated port. (Default: Disabled)
- ◆ **MAC Filter ID** – Allows a MAC Filter to be assigned to the port. MAC addresses or MAC address ranges present in a selected MAC Filter are exempt from authentication on the specified port (as described under "[Configuring a MAC Address Filter](#)"). (Range: 1-64; Default: None)

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 195: Configuring Interface Settings for Network Access

Security > Network Access

Step: 2. Configure Interface

☒ General ☐ Link Detection

Port List Total: 38

Port	Status	MAC Authentication		Network Access	Guest VLAN	Dynamic VLAN	Dynamic QoS	MAC Filter ID
		Intrusion	Max MAC Count (1-1024)	Max MAC Count (1-2048)	(1-4094, 0: Disabled)			(1-64)
1	☐ Enabled	Block	1024	1024	0	☒ Enabled	☐ Enabled	
2	☐ Enabled	Block	1024	1024	0	☒ Enabled	☐ Enabled	
3	☐ Enabled	Block	1024	1024	0	☒ Enabled	☐ Enabled	
4	☐ Enabled	Block	1024	1024	0	☒ Enabled	☐ Enabled	
5	☐ Enabled	Block	1024	1024	0	☒ Enabled	☐ Enabled	

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 196: Configuring Link Detection for Network Access

Security > Network Access

Step: 2. Configure Interface

General Link Detection

Port List Total: 20

Port	Link Detection Status	Condition	Action
1	☐ Enabled	Link down	Trap
2	☐ Enabled	Link 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 197: 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 (XX-XX-XX-XX-XX-XX OF 000000000000)

MAC Address Mask FFFFFFFF (XX-XX-XX-XX-XX-XX OF 000000000000)

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 198: Showing the MAC Address Filter Table for Network Access

Security > Network Access

Step: 3. Configure MAC Filter Action: Show

MAC Filter List Total: 1

	Filter ID	MAC Address	MAC Address Mask
☐	22	11-22-33-44-55-66	FF-FF-FF-FF-FF-FF

Delete 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 199: Showing Addresses Authenticated for Network Access

Security > Network Access

Step: 4. Show Information

Query By:

Sort Key: (00:00:00:00:00:00)

☐ MAC Address (00:00:00:00:00:00)

☐ Interface

☐ Attribute

Authenticated MAC Address List Total: 6

☐	MAC Address	Interface	RADIUS Server	Time	Attribute
☐	00-30-06-45-F2-23	Unit 1 / Port 23	10.2.2.10	2008y 26m 12d 11h 16m 12s	Dynamic
☐	00-68-E8-5E-E1-00	Unit 1 / Port 23	10.2.2.10	2008y 26m 12d 11h 32m 24s	Dynamic
☐	00-30-06-45-F2-23	Unit 1 / Port 23	10.2.2.10	2008y 26m 12d 11h 40m 32s	Dynamic
☐	00-41-88-31-B8-30	Unit 1 / Port 23	10.2.2.10	2008y 26m 12d 11h 18m 51s	Dynamic
☐	00-41-88-36-05-28	Unit 1 / Port 23	10.2.2.10	2008y 26m 12d 11h 32m 22s	Dynamic
☐	00-01-83-36-D3-7F	Unit 1 / Port 23	10.2.2.10	2008y 26m 12d 11h 32m 28s	Dynamic
☐	00-01-83-3C-3C-18	Unit 2 / Port 23	10.2.2.10	2008y 26m 12d 11h 15m 19s	Dynamic
☐	00-01-83-3C-3E-82	Unit 2 / Port 23	10.2.2.10	2008y 26m 12d 11h 17m 48s	Dynamic

Configuring HTTPS

You can configure the switch to enable the Secure Hypertext Transfer Protocol (HTTPS) over the Secure Socket Layer (SSL), providing secure access (i.e., an encrypted connection) to the switch's web interface.

Configuring Global Settings for HTTPS

Use the Security > HTTPS (Configure Global) page to enable or disable HTTPS and specify the TCP port used for this service.

Command Usage

- ◆ Both the HTTP and HTTPS service can be enabled independently on the switch. However, you cannot configure both services to use the same TCP port. (HTTP can only be configured through the CLI using the "ip http server" command described in the *CLI Reference Guide*.)
- ◆ If you enable HTTPS, you must indicate this in the URL that you specify in your browser: `https://device[:port_number]`
- ◆ When you start HTTPS, the connection is established in this way:
 - The client authenticates the server using the server's digital certificate.
 - The client and server negotiate a set of security protocols to use for the connection.

- The client and server generate session keys for encrypting and decrypting data.
- ◆ The client and server establish a secure encrypted connection.
A padlock icon should appear in the status bar for Internet Explorer 9, Mozilla Firefox 52, Google Chrome 54, or Opera 41 or more recent versions.
- ◆ The following web browsers and operating systems currently support HTTPS:

Table 21: HTTPS System Support

Web Browser	Operating System
Internet Explorer 9.x or later	Windows 7, 8, 10
Mozilla Firefox 52 or later	Windows 7, 8, 10, Linux
Google Chrome 54 or later	Windows 7, 8, 10
Opera 41 or later	Windows 7, 8, 10, Linux

- ◆ To specify a secure-site certificate, see [“Replacing the Default Secure-site Certificate” on page 319](#).



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.

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 200: Configuring HTTPS

The screenshot shows a web interface for configuring HTTPS. At the top, it says 'Security > HTTPS'. Below that is an 'Actions:' dropdown menu currently set to 'Configure Global'. Underneath, there are two settings: 'HTTPS Status' with a checked checkbox and the text 'enabled', and 'UDP Port (1-65535)' with a text box containing '443'. At the bottom right, there are two buttons: 'Apply' and '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 103](#) 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.
- ◆ **Delete** – Deletes the HTTPS secure site certificate. You must reboot the switch to load the default certificate.

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 201: Downloading the Secure-Site Certificate

The screenshot shows a web interface for configuring HTTPS. The title bar reads "Security > HTTPS". Below the title, there is a dropdown menu labeled "Action:" with "Copy Certificate" selected. The main form contains several input fields: "TFTP Server IP Address" (192.168.0.4), "Certificate Source File Name" (site-certificate), "Private Key Source File Name" (private-key), "Private Password" (password), and "Confirm Password" (password). At the bottom right, there are "Apply" and "Revert" buttons. At the bottom left, there is a "Delete" button with a tooltip that says "Click this button to delete current certificate."

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 only SSH Version 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 287](#)). 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 325](#) 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 303](#).) 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 key:

```
-----BEGIN RSA PUBLIC KEY-----
MIIBCgKCAQEAmk9QvzWeYFgvjInC/jyvalribLTmx5ncPJcnC0RHsAgbuzefGRhV
RSdcx63mRUXiG968LksVXPkJEa6cX02xZoLwbg96VpoarLTrnI8UC55XFw4c32I
q0FERXUIVQPfSt2XgPqOH+uIaokeKJ6SV3V05CmdXwRYmSr4ZGQB2SjOp4cMpmY
TKIIfHxQ0L7WfSKJnBMgsk9ZyzvDlpZxXoHKdG87k4k9gn8d+fcV9d7xHf2255jI
1xNwPOZDH57yjc63s04xoYcTE7TBrCM5vqJwr17R8ioNXorpRvWZ1ump2FZTZHE
nyuoYvSJMaR8iTTboFIjEonqVX6gsLFwhwIDAQAB
-----END RSA PUBLIC KEY-----
```

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 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 v2 Clients

- a. The client first queries the switch to determine if 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 RSA host keys before enabling the SSH server. See [“Generating the Host Key Pair” on page 324](#).

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 and the switch supports management access only from SSH Version 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)

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 202: Configuring the SSH Server



Security > SSH

Step: 1. Configure Global

SSH Server Status: ☐ Enabled

Version: 2.0

Authentication Timeout (1-120): 120 sec

Authentication Retries (1-5): 3

Apply Revert

Generating the Host Key Pair

Use the Security > SSH (Configure Host Key - Generate) page to generate a host public/private key pair used to provide secure communications between an SSH client and the switch. After generating this key pair, you must provide the host public key to SSH clients and import the client's public key to the switch as described in the section ["Importing User Public Keys" on page 325](#).



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 323](#).

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)
The SSH server uses RSA 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 keys for SSHv2 clients.

Web Interface

To generate the SSH host key pair:

1. Click Security, SSH.
2. Select Configure Host Key from the Step list.
3. Select Generate from the Action list.
4. Click Apply.

Figure 203: Generating the SSH Host Key Pair

Security > SSH

Step: 2. Configure Host Key Action: Generate

Host-Key Type: RSA

Apply Revert

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 204: Showing the SSH Host Key Pair

Security > SSH

Step: 3. Configure Host Key Action: Show

Public-Key of Host-Key

-----BEGIN RSA PUBLIC KEY-----
MIIBBgMCAgEAAQ/IMY2iNyNDsBeuSepFCotjahnZWF1S9/GIDwyFTHtGFEStMxU1z
+STRN90ymKld3nsFjBzRvEXILE82IS8uQkQPm0cn3FNI1HI/JPvIksqF1wu5uw
FCma3Arcn3lGcr3SEGbJZFAG6ZVeJcSS2wTQQ5UeSeajRFvaKScspJtq/EeYtJ
H1IF+qABRm06R/vlc3HUQxeDtwF1yphT6GHyatwQqdatnRgkags4IpozmkJPlqqs

Clear

Save Click this button to Save All Host-Key from Memory to Flash

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 303](#)).
- ◆ **User Key Type** – The type of public key to upload.
 - RSA: The switch accepts an RSA encrypted public key.

The SSH server uses RSA 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 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 from the drop-down box, input the TFTP server IP address and the public key source file name.
5. Click Apply.

Figure 205: Copying the SSH User's Public Key

The screenshot shows the 'Security > SSH' configuration page. At the top, there are two dropdown menus: 'Step: 3. Configure User Key' and 'Action: Copy'. Below these, there are four input fields: 'User Name' (a dropdown menu with 'admin' selected), 'User-Key Type' (a text field with 'RSA' entered), 'TFTP Server IP Address' (an empty text field), and 'Source File Name' (an empty text field). At the bottom right of the form, there are two buttons: 'Apply' and 'Revert'.

To display or clear the SSH user's public key:

1. Click Security, SSH.
2. Select Configure User Key from the Step list.
3. Select Show from the Action list.
4. Select a user from the User Name list.
5. Select the host-key type to clear.
6. Click Clear.

Figure 206: Showing the SSH User's Public Key

Security > SSH

Step: 1. Configure User Key Action: Show

User Name: admin

Public Key of User Key

☒ RSA

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 64.
- ◆ The maximum number of rules per system is 1024 rules.
- ◆ An ACL can have up to 1024 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.

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:


- ◆ **Pool Capability Code** – Abbreviation for processes shown in the TCAM List.
- ◆ **Unit** – Stack unit identifier.
- ◆ **Device** – Memory chip used for indicated pools.
- ◆ **Pool** – Rule slice (or call group). Each slice has a fixed number of rules that are used for the specified features.
- ◆ **Total** – The maximum number of policy control entries allocated to the each pool.
- ◆ **Used** – The number of policy control entries used by the operating system.
- ◆ **Free** – The number of policy control entries available for use.
- ◆ **Capability** – The processes assigned to each pool.

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 207: Showing TCAM Utilization



Setting the ACL Name and Type

Use the Security > ACL (Configure ACL - Add) page to create an ACL.

Parameters

These parameters are displayed:

- ◆ **ACL Name** – Name of the ACL. (Maximum length: 32 characters)
- ◆ **Type** – The following filter modes are supported:
 - **IP Standard:** IPv4 ACL mode filters packets based on the source IPv4 address.
 - **IP Extended:** IPv4 ACL mode filters packets based on the source or destination IPv4 address, as well as the protocol type and protocol port number. If the “TCP” protocol is specified, then you can also filter packets based on the TCP control code.
 - **IPv6 Standard:** IPv6 ACL mode filters packets based on the source IPv6 address.
 - **IPv6 Extended:** IPv6 ACL mode filters packets based on the source or destination IP address, 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 399](#)).

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 208: Creating an ACL

Security > ACL

Step: 1. 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 209: Showing a List of ACLs

Security > ACL

Step: 1. Configure ACL Action: Show

ACL List Total: 10

	ACL Name	Type
☐	acStandard1	IP Standard
☐	acStandard2	IP Standard
☐	acExtended1	IP Extended
☐	acExtended2	IP Extended
☐	acMAC1	MAC
☐	acMAC2	MAC
☐	acMAC3	MAC
☐	acIPv6Standard	IPv6 Standard
☐	acIPv6Extended	IPv6 Extended
☐	acADP	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 210: Configuring a Standard IPv4 ACL

The screenshot shows the 'Security > ACL' configuration page. At the top, the 'Step' is '1. Configure ACL' and the 'Action' is 'Add Rule'. Below this, there are radio buttons for 'Type': 'IP Standard' (selected), 'IP Extended', 'MAC', 'IPv6 Standard', 'IPv6 Extended', and 'ARP'. The 'Name' field is set to 'R&D'. The 'Action' dropdown is set to 'Permit'. The 'Address Type' dropdown is set to 'Host'. The 'Source IP Address' field is '10.1.1.21' and the 'Source Subnet Mask' field is '255.255.255.255'. The 'Time-Range' checkbox is checked, and its dropdown is set to 'R&D'. At the bottom right 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 332](#).)
- ◆ **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 ICMP, TCP, UDP or Others, where others indicates a specific protocol number (0-255). (Options: ICMP, TCP, UDP, Others; Default: Others)

The following items are under TCP

- **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

◆ **Service Type** – Packet priority settings based on the following criteria:

- **Precedence** – IP precedence level. (Range: 0-7)
- **DSCP** – DSCP priority level. (Range: 0-63)

◆ **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 211: Configuring an Extended IPv4 ACL

The screenshot shows the 'Security > ACL' configuration page. At the top, there's a breadcrumb 'Security > ACL' and a green plus icon. Below it, a 'Step' dropdown is set to '1. Configure ACL' and an 'Action' dropdown is set to 'Add Rule'. The 'Type' section has radio buttons for 'IP Standard', 'IP Extended' (selected), 'MAC', 'IPv6 Standard', 'IPv6 Extended', and 'ARP'. Below 'Type' is a 'Name' dropdown. The 'Action' section has a 'Permit' dropdown. The 'Source Address Type' is 'Any' and the 'Destination Address Type' is 'Any'. Both have corresponding 'Source IP Address' and 'Destination IP Address' fields set to '0.0.0.0'. Below these are 'Source Subnet Mask' and 'Destination Subnet Mask' fields, both set to '0.0.0.0'. There are also fields for 'Source Port (0-65535)', 'Destination Port (0-65535)', 'Source Port Bit Mask (0-65535)', and 'Destination Port Bit Mask (0-65535)'. The 'Protocol' section has radio buttons for 'ICMP (1)', 'TCP (6)', 'UDP (17)', and 'Others' (selected). Below 'Others' are fields for 'ICMP Type (0-255)', 'Control Code (0-63)', and 'Control Code Bit Mask (0-63)'. The 'Service Type' section has radio buttons for 'Precedence (0-7)' (selected) and 'DSCP (0-63)'. Below 'Precedence (0-7)' is a field. At the bottom left is a 'Time-Range' dropdown set to 'Any'. At the bottom right are 'Apply' and 'Reset' buttons.

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 212: Configuring a Standard IPv6 ACL

The screenshot shows the 'Security > ACL' configuration page. The 'Step' is '1. Configure ACL' and the 'Action' is 'Add Rule'. Under 'Type', 'IPv6 Standard' is selected. The 'Name' is 'R&D@BS'. Under 'Action', 'Permit' is selected. Under 'Source Address Type', 'Host' is selected. The 'Source IPv6 Address' is '2001:DB8:2229::79'. The 'Source Prefix Length (0-128)' is '128'. The 'Time-Range' checkbox is unchecked. 'Apply' and 'Revert' buttons are at the bottom right.

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.
- ◆ **Protocol** – Selects the protocol of the next header in the packet. Select TCP, UDP, ICMP, or Next Header to identify the protocol by value.
- ◆ **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)

- ◆ **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)
- ◆ **Source Port** – Protocol⁹ source port number. (Range: 0-65535)
- ◆ **Source Port Bit Mask** – Decimal number representing the port bits to match. (Range: 0-65535)
- ◆ **Destination Port** – Protocol⁹ destination port number. (Range: 0-65535)
- ◆ **Destination Port Bit Mask** – Decimal number representing the port bits to match. (Range: 0-65535)
- ◆ **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).

9. Includes TCP, UDP or other protocol types.

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 213: Configuring an Extended IPv6 ACL

The screenshot shows a configuration window for an Extended IPv6 ACL. The 'Action' is set to 'Permit'. The 'Protocol' is checked. The 'Next-Header' is set to 'Next-Header'. The 'Source Address Type' is set to 'Any'. The 'Source IPv6 Address' is set to '::'. The 'Source Prefix Length' is set to '0'. The 'Destination Address Type' is set to 'Any'. The 'Destination IPv6 Address' is set to '::'. The 'Destination Prefix Length' is set to '0'. The 'DSCP' is set to '0'. The 'Time-Range' is set to 'Any'. The 'Apply' and 'Reset' 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 333.](#)
- **IPv6** – [See “Configuring an Extended IPv6 ACL” on page 337.](#)

◆ **CoS** – CoS value. (Range: 0-7, where 7 is the highest priority)

◆ **CoS Bit Mask** – CoS bitmask. (Range: 0-7)

◆ **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 214: Configuring a MAC ACL

The screenshot shows the 'Configure ACL - Add Rule' configuration page. The 'Step' is '1 Configure ACL' and the 'Action' is 'Add Rule'. The 'Type' is set to 'MAC'. The 'Name' field is empty. The 'Action' is set to 'Permit'. The 'Source Address Type' is 'Any'. The 'Source MAC Address' is '00-00-00-00-00-00'. The 'Source Bit Mask' is '00-00-00-00-00-00'. The 'Destination Address Type' is 'Any'. The 'Destination MAC Address' is '00-00-00-00-00-00'. The 'Destination Bit Mask' is '00-00-00-00-00-00'. The 'Packet Format' is 'Any'. The 'VID (1-4094)' field is empty. The 'VID Bit Mask (0-4095)' field is empty. The 'Ethernet Type' field is empty, with a hint '(0000-FFFF, hexadecimal value)'. The 'Ethernet Type Bit Mask' field is empty, with a hint '(0000-FFFF, hexadecimal value)'. The 'Internet Protocol' is set to 'No'. The 'Cos (0-7)' field is empty. The 'Cos Bit Mask (0-7)' field is empty. The 'Time-Range' field is empty. The 'Apply' and 'Revert' buttons are at the bottom right.

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 400](#)).

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 332](#).)
- ◆ **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 215: Configuring a ARP ACL

The screenshot shows the 'Security > ACL' configuration page. The 'Step' dropdown is set to '1. Configure ACL' and the 'Action' dropdown is set to 'Add Rule'. Under the 'Type' section, 'ARP' is selected. The 'Name' field contains 'R8F8745P'. The 'Action' is set to 'Permit'. The 'Packet Type' is set to 'P'. The 'Source IP Address Type' is set to 'Any', and the 'Destination IP Address Type' is set to 'Any'. The 'Source IP Address' and 'Destination IP Address' fields are empty. The 'Source IP Subnet Mask' and 'Destination IP Subnet Mask' fields are empty. The 'Source MAC Address Type' is set to 'Any', and the 'Destination MAC Address Type' is set to 'Any'. The 'Source MAC Address' and 'Destination MAC Address' fields are empty. The 'Source MAC Bit Mask' and 'Destination MAC Bit Mask' fields are empty. There is a 'Log' checkbox and 'Apply' and 'Revert' buttons at the bottom.

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.

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 216: Binding a Port to an ACL

The screenshot shows a web-based configuration interface for ACLs. The breadcrumb is 'Security > ACL'. The 'Step' is '2. Configure Interface' and the 'Action' is 'Configure'. There are three radio buttons for 'Type': IP (selected), MAC, and IPv6. Below this is a 'Port' dropdown menu set to '1'. Under the 'Ingress' section, there is an 'ACL' dropdown set to '1', a 'Time-Range' dropdown, and a 'Counter' checkbox. A similar set of options is provided for the 'Egress' section. At the bottom right, there are 'Apply' and 'Revert' buttons.

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: 1-28/54)
- ◆ **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.
- ◆ **Clear Counter** – Clears the hit counter for the specified ACL.

Web Interface

To show statistics for ACL hardware counters:

1. Click Security, ACL.
2. Select Configure Interface from the Step list.
3. Select Show Hardware Counters from the Action list.
4. Select a port.
5. Select ingress or egress traffic.

Figure 217: Showing ACL Statistics

The screenshot shows the 'Security > ACL' web interface. At the top, there's a breadcrumb 'Security > ACL' and a green help icon. Below it, a 'Step' dropdown is set to '2. Configure Interface' and an 'Action' dropdown is set to 'Show Hardware Counter'. The main configuration area has three dropdowns: 'Port' set to '1', 'Type' set to 'IP Standard', and 'Direction' set to 'Ingress'. A 'Query' button is to the right. Below the configuration area, the 'ACL Name' is 'no'. A 'Total' label is followed by a table with 5 columns: 'Action', 'Source IP Address', 'Time Range', 'Hit', and 'Clear Counter'. The table has one row with 'Permit', 'Any', an empty 'Time Range' cell, '14', and a 'Clear' button.

Action	Source IP Address	Time Range	Hit	Clear Counter
Permit	Any		14	Clear

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 218: Creating an IP Address Filter for Management Access

The screenshot shows the 'Security > IP Filter' web interface. At the top, there is a breadcrumb 'Security > IP Filter'. Below it, the 'Action' dropdown is set to 'Add'. The 'Mode' section has 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 empty. At the bottom right, there are 'Apply' and 'Revert' buttons.

To show a list of IP addresses authorized for management access:

1. Click Security, IP Filter.
2. Select Show from the Action list.

Figure 219: Showing IP Addresses Authorized for Management Access

The screenshot shows the 'Security > IP Filter' web interface with the 'Action' dropdown set to 'Show'. The 'Mode' section has four radio buttons: 'Web', 'SNMP' (selected), 'Telnet', and 'All'. Below the mode selection, there is a table titled 'SNMP IP Filter List Table 1'. The table has two columns: 'Start IP Address' and 'End IP Address'. There is one row in the table with the values '10.1.2.3' and '10.1.2.3'. At the bottom right, there are 'Delete' and 'Revert' buttons.

Start IP Address	End IP Address
10.1.2.3	10.1.2.3

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 205](#)).

- ◆ 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 108](#)).
- ◆ 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 314](#).
- ◆ **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.
- ◆ **Sticky MAC** – Sticky MAC addresses that port security has learned are dynamic addresses that cannot be moved to another port. If sticky MAC addresses are received on another secure port, then the port intrusion action is taken.

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 220: Configuring Port Security



The screenshot shows the 'Security > Port Security' configuration page. It features a 'Port Security List' table with 11 columns: Port, Security Status, Port Status, Action, Max MAC Count (0-9999), Current MAC Count, MAC Filter, MAC Filter ID, Last Intrusion MAC, Last Time Detected Intrusion MAC, and Sticky MAC. There are four rows representing ports 1 through 4. Each row has a checkbox for 'Security Status' (all are checked), a 'Port Status' dropdown (all are 'Secure/Down'), an 'Action' dropdown (all are 'None'), a 'Max MAC Count' input field (all are '0'), a 'Current MAC Count' input field (all are '0'), a 'MAC Filter' dropdown (all are 'Disabled'), a 'MAC Filter ID' input field (all are '0'), a 'Last Intrusion MAC' input field (all are 'NA'), a 'Last Time Detected Intrusion MAC' input field (all are 'NA'), and a 'Sticky MAC' checkbox (all are checked).

Port	Security Status	Port Status	Action	Max MAC Count (0-9999)	Current MAC Count	MAC Filter	MAC Filter ID	Last Intrusion MAC	Last Time Detected Intrusion MAC	Sticky MAC
1	☒	Secure/Down	None	0	0	Disabled	0	NA	NA	☒
2	☒	Secure/Down	None	0	0	Disabled	0	NA	NA	☒
3	☒	Secure/Down	None	0	0	Disabled	0	NA	NA	☒
4	☒	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 221: 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 7, 8 and 10.

Configuring 802.1X Global Settings Use the Security > Port Authentication (Configure Global) page to configure IEEE 802.1X port authentication. The 802.1X protocol must be enabled globally for the switch system before port settings are active.

Parameters

These parameters are displayed:

- ◆ **System Authentication Control** – Sets the global setting for 802.1X. (Default: Disabled)
- ◆ **EAPOL Pass Through** – Passes EAPOL frames through to all ports in STP forwarding state when dot1x is globally disabled. (Default: Disabled)

When this device is functioning as intermediate node in the network and does not need to perform dot1x authentication, **EAPOL Pass Through** can be enabled to allow the switch to forward EAPOL frames from other switches on to the authentication servers, thereby allowing the authentication process to still be carried out by switches located on the edge of the network.

When this device is functioning as an edge switch but does not require any attached clients to be authenticated, **EAPOL Pass Through** can be disabled to discard unnecessary EAPOL traffic.
- ◆ **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 357](#)).
- ◆ **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.
4. Click Apply

Figure 222: Configuring Global Settings for 802.1X Port Authentication

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.
- ◆ 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 288.](#))
- ◆ **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 165](#)) and mapped on each port (See [“Configuring Network Access for Ports” on page 311](#)).

Supplicant List

- ◆ **Supplicant** – MAC address of authorized client.

Authenticator PAE State Machine

- ◆ **State** – Current state (including initialize, disconnected, connecting, authenticating, authenticated, aborting, held, force_authorized, force_unauthorized).
- ◆ **Reauth Count** – Number of times connecting state is re-entered.
- ◆ **Current Identifier** – Identifier sent in each EAP Success, Failure or Request packet by the Authentication Server.

Backend State Machine

- ◆ **State** – Current state (including request, response, success, fail, timeout, idle, initialize).
- ◆ **Request Count** – Number of EAP Request packets sent to the Supplicant without receiving a response.
- ◆ **Identifier (Server)** – Identifier carried in the most recent EAP Success, Failure or Request packet received from the Authentication Server.

Reauthentication State Machine

- ◆ **State** – Current state (including initialize, reauthenticate).

Web Interface

To configure port authenticator settings for 802.1X:

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

Figure 223: 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): 1

Max Request (1-10): 2

Quiet Period (1-65535): 60 sec

Tx Period (1-65535): 20 sec

Supplicant Timeout (1-65535): 30 sec

Server Timeout: 0 sec

Re-authentication Status: ☐ Enabled

Re-authentication Period (1-65535): 3000 sec

Re-authentication Max Retries (1-10): 2

Intrusion Action: Block Traffic

Supplicant List: Total: 0

Supplicant	Authenticator PAE Slave Machine		Backend Slave Machine		Reauthentication Slave Machine	
	State	Reauth Count	Current Identifier	State	Request Count	Identifier (Server)

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 353](#)), 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 352](#)) 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 353](#)).
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 224: Configuring Interface Settings for 802.1X Port Supplicant

The screenshot shows the 'Security > Port Authentication' configuration page. The 'Step' dropdown is set to '2. Configure Interface'. Under 'Type', 'Authenticator' is unselected and 'Supplicant' is selected. The 'Port' is set to '1'. The 'PAE Supplicant' checkbox is checked, with the label 'Enabled' next to it. Below this are five input fields: 'Authentication Period (1-65530)' with value '30', 'Hold Period (1-65530)' with value '60', 'Start Period (1-65530)' with value '30', 'Maximum Start (1-65530)' with value '3', and 'Authenticated' with value 'No'. At the bottom right are 'Apply' and 'Revert' buttons.

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.

Table 22: 802.1X Statistics (Continued)


Parameter	Description
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.
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.
- 4. Select a port.


Figure 225: Showing Statistics for 802.1X Port Authenticator



To display port supplicant statistics for 802.1X:

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

Figure 226: Showing Statistics for 802.1X Port Supplicant



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: Disabled)
- ◆ **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: Disabled)
- ◆ **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: Disabled)
- ◆ **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: Disabled)
- ◆ **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: Disabled)
- ◆ **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 227: Protecting Against DoS Attacks

The screenshot shows a web interface titled "Security > DoS Protection". It contains a list of DoS attack types with checkboxes for enabling protection and input fields for setting the maximum allowed rate in kbps. The settings are as follows:

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

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

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 Security > 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 Remote ID TR101 VLAN Field** – Adds ":VLAN" in TR101 field for untagged packets.

The format for TR101 option 82 is: "<IP> eth <SID>/<PORT>[:<VLAN>]". Note that the SID (Switch ID) is always 0. By default the PVID is added to the end of the TR101 field for untagged packets. For tagged packets, the VLAN ID is always added.

- ◆ **DHCP Snooping Information Option TR101 Board ID** – Sets the board identifier used in Option 82 information based on TR-101 syntax. (Range: 0-9; Default: undefined)
- ◆ **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 228: Configuring Global Settings for DHCP Snooping

The screenshot shows the 'Security > DHCP Snooping' configuration page. At the top, there's a breadcrumb trail and a 'Step: 1. Configure Global' dropdown. The page is divided into two main sections: 'General' and 'Information'.
General Section:
- 'DHCP Snooping Status' has a checkbox for 'Enabled' which is currently unchecked.
- 'DHCP Snooping MAC-Address Verification' has a checkbox for 'Enabled' which is checked.
- 'DHCP Snooping Rate Limit' has a checkbox for 'Enabled' which is unchecked, followed by a text input field containing '0' and the unit '(pkts/sec)'.
Information Section:
- 'DHCP Snooping Information Option Status' has a checkbox for 'Enabled' which is unchecked.
- 'DHCP Snooping Information Option Sub-option Format' has a dropdown menu set to 'Extra Subtype Included'.
- 'DHCP Snooping Information Option Remote ID' has a dropdown menu set to 'MAC Address (Hex Encoded)' followed by a text input field.
- 'DHCP Snooping Information Option Remote ID TR101 VLAN Field' has a checkbox for 'Enabled' which is checked.
- 'DHCP Snooping Information Option TR101 Board ID' has a text input field containing 'none'.
- 'DHCP Snooping Information Option Policy' has a dropdown menu set to 'Replace'.
At the bottom right, there are 'Apply' and 'Revert' buttons.

DHCP Snooping VLAN Configuration Use the Security > 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 229: Configuring DHCP Snooping on a VLAN

The screenshot shows a web interface for configuring DHCP Snooping. The breadcrumb trail at the top is "IP Service > DHCP > Snooping". Below this, there is a "Step:" dropdown menu currently showing "2: Configure VLAN". The main configuration area contains two fields: "VLAN" with a value of "1" and a dropdown arrow, and "DHCP Snooping Status" with a checked checkbox and the text "Enabled". At the bottom right of the form are two buttons: "Apply" and "Revert".

Configuring Ports for DHCP Snooping Use the Security > 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.
- ◆ The format for TR101 option 82 is: "<IP> eth <SID>/<PORT>[:<VLAN>]". Note that the SID (Switch ID) is always 0. By default the PVID is added to the end of the TR101 field for untagged packets. For tagged packets, the VLAN ID is always added.

Parameters

These parameters are displayed:

- ◆ **Trust Status** – Enables or disables a port as trusted. (Default: Disabled)
- ◆ **Max Number** – The maximum number of DHCP clients which can be supported per interface. (Range: 1-32; Default: 16)
- ◆ **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)
 - **TR101 VLAN Field** – Adds ":VLAN" in TR101 field for untagged packets.

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 230: Configuring the Port Mode for DHCP Snooping

Port	Trust Status	Max Number (1-10)	Mode	Circuit ID	TR101 VLAN Field
1	☒ Enabled	10	VLAN-User-Port		☒ Enabled
2	☒ Enabled	10	VLAN-User-Port		☒ Enabled
3	☒ Enabled	10	VLAN-User-Port		☒ Enabled
4	☒ Enabled	10	VLAN-User-Port		☒ Enabled
5	☒ Enabled	10	VLAN-User-Port		☒ Enabled

Displaying DHCP Snooping Binding Information

Use the Security > 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-DHCPSPN** – 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 231: Displaying the Binding Table for DHCP Snooping



IP Service > DHCP > Snooping

Step: 4. Show Information

DHCP Snooping Binding List Total: 6

MAC Address	IP Address	Lease Time (seconds)	Type	VLAN	Interface
10-10-85-74-00-01	10.2.44.90	5	DHCP-Snooping	1	Trunk 1
10-10-85-74-00-02	10.3.44.90	15	Static-DHCP-SNP	1	Int 1 / Port 2
10-10-85-74-00-03	10.4.44.90	25	DHCP-Snooping	1	Int 1 / Port 3
10-10-85-74-00-04	10.5.44.90	15	Static-DHCP-SNP	1	Trunk 4
10-10-85-74-00-05	10.6.44.90	15	DHCP-Snooping	1	Int 1 / Port 5
10-10-85-74-00-06	10.7.44.90	5	Static-DHCP-SNP	1	Int 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.

DHCPv6 Snooping

The addresses assigned to DHCPv6 clients on insecure ports can be carefully controlled using the dynamic bindings registered with DHCPv6 Snooping (or using the static bindings configured with IPv6 Source Guard). DHCPv6 snooping allows a switch to protect a network from rogue DHCPv6 servers or other devices which send port-related information to a DHCPv6 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 DHCPv6 messages are received from an outside source. DHCPv6 snooping is used to filter DHCPv6 messages received on a unsecured interface from outside the network or fire wall. When DHCPv6 snooping is enabled globally and enabled on a VLAN interface, DHCPv6 messages received on an untrusted interface from a device not listed in the DHCPv6 snooping table will be dropped.
- ◆ When enabled, DHCPv6 messages entering an untrusted interface are filtered based upon dynamic entries learned via DHCPv6 snooping.

- ◆ Table entries are only learned for trusted interfaces. Each entry includes a MAC address, IPv6 address, lease time, binding type, VLAN identifier, and port identifier.
- ◆ When DHCPv6 snooping is enabled, the rate limit for the number of DHCPv6 messages that can be processed by the switch is 100 packets per second. Any DHCPv6 packets in excess of this limit are dropped.
- ◆ Filtering rules are implemented as follows:
 - If global DHCPv6 snooping is disabled, all DHCPv6 packets are forwarded.
 - If DHCPv6 snooping is enabled globally, and also enabled on the VLAN where the DHCPv6 packet is received, DHCPv6 packets are forwarded for a *trusted* port as described below.
 - If DHCPv6 snooping is enabled globally, and also enabled on the VLAN where the DHCP packet is received, but the port is *not trusted*, DHCP packets are processed according to message type as follows:

DHCP Client Packet

- Request: Update entry in binding cache, recording client's DHCPv6 Unique Identifier (DUID), server's DUID, Identity Association (IA) type, IA Identifier, and address (4 message exchanges to get IPv6 address), and forward to trusted port.
- Solicit: Add new entry in binding cache, recording client's DUID, IA type, IA ID (2 message exchanges to get IPv6 address with rapid commit option, otherwise 4 message exchanges), and forward to trusted port.
- Decline: If no matching entry is found in binding cache, drop this packet.
- Renew, Rebind, Release, Confirm: If no matching entry is found in binding cache, drop this packet.
- If the DHCPv6 packet is not a recognizable type, it is dropped.

If a DHCPv6 packet from a client passes the filtering criteria above, it will only be forwarded to trusted ports in the same VLAN.

DHCP Server Packet

- If a DHCP server packet is received on an *untrusted* port, drop this packet and add a log entry in the system.
- If a DHCPv6 Reply packet is received from a server on a *trusted* port, it will be processed in the following manner:

A. Check if IPv6 address in IA option is found in binding table:

- If yes, continue to C.
 - If not, continue to B.
 - Check if IPv6 address in IA option is found in binding cache:
 - If yes, continue to C.
 - If not, check failed, and forward packet to trusted port.
- B.** Check status code in IA option:
- If successful, and entry is in binding table, update lease time and forward to original destination.
 - If successful, and entry is in binding cache, move entry from binding cache to binding table, update lease time and forward to original destination.
 - Otherwise, remove binding entry. and check failed.
 - If a DHCPv6 Relay packet is received, check the relay message option in Relay-Forward or Relay-Reply packet, and process client and server packets as described above.
- ◆ If DHCPv6 snooping is globally disabled, all dynamic bindings are removed from the binding table.
 - ◆ *Additional considerations when the switch itself is a DHCPv6 client* – The port(s) through which the switch submits a client request to the DHCPv6 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 DHCPv6 server. Also, when the switch sends out DHCPv6 client packets for itself, no filtering takes place. However, when the switch receives any messages from a DHCPv6 server, any packets received from untrusted ports are dropped.

DHCPv6 Snooping Global Configuration Use the Security > DHCP Snooping6 (Configure Global) page to enable DHCPv6 Snooping globally on the switch, or to configure MAC Address Verification.

Parameters

These parameters are displayed:

- ◆ **DHCPv6 Snooping Status** – Enables DHCPv6 snooping globally. (Default: Disabled)
- ◆ **DHCPv6 Snooping Option Remote ID** – Enables the insertion of remote-id option 37 information into DHCPv6 client messages. Remote-id option information such as the port attached to the client, DUID, and VLAN ID is used

by the DHCPv6 server to assign preassigned configuration data specific to the DHCPv6 client. (Default: Disabled)

- DHCPv6 provides a relay mechanism for sending information about the switch and its DHCPv6 clients to the DHCPv6 server. Known as DHCPv6 Option 37, it allows compatible DHCPv6 servers to use the information when assigning IP addresses, or to set other services or policies for clients.
 - When DHCPv6 Snooping Information Option 37 is enabled, the requesting client (or an intermediate relay agent that has used the information fields to describe itself) can be identified in the DHCPv6 request packets forwarded by the switch and in reply packets sent back from the DHCPv6 server.
 - When the DHCPv6 Snooping Option 37 is enabled, clients can be identified by the switch port to which they are connected rather than just their MAC address. DHCPv6 client-server exchange messages are then forwarded directly between the server and client without having to flood them to the entire VLAN.
 - DHCPv6 snooping must be enabled for the DHCPv6 Option 37 information to be inserted into packets. When enabled, the switch will either drop, keep or remove option 37 information in incoming DHCPv6 packets. Packets are processed as follows:
 - If an incoming packet is a DHCPv6 request packet with option 37 information, it will modify the option 37 information according to the settings specified.
 - If an incoming packet is a DHCPv6 request packet without option 37 information, enabling the DHCPv6 snooping information option will add option 37 information to the packet.
 - If an incoming packet is a DHCPv6 reply packet with option 37 information, enabling the DHCPv6 snooping information option will remove option 37 information from the packet.
 - When this switch inserts Option 37 information in DHCPv6 client request packets, the switch's MAC address (hexadecimal) is used for the remote ID.
- ◆ **DHCPv6 Snooping Option Policy** – Sets the remote-id option policy for DHCPv6 client packets that include Option 37 information.

When the switch receives DHCPv6 packets from clients that already include DHCP Option 37 information, the switch can be configured to set the action policy for these packets. The switch can either drop the DHCPv6 packets, keep the existing information, or replace it with the switch's relay agent information.

- **Drop** – Drops the client's request packet instead of relaying it. (This is the default policy.)
- **Keep** – Retains the Option 82 information in the client request, and forwards the packets to trusted ports.

- **Replace** – Replaces the Option 37 remote-ID in the client’s request with the relay agent’s remote-ID (when DHCPv6 snooping is enabled), and forwards the packets to trusted ports.

Web Interface

To configure global settings for DHCPv6 Snooping:

1. Click Security, DHCP Snooping6.
2. Select Configure Global from the Step list.
3. Select the required options for the DHCPv6 snooping process and for the DHCPv6 snooping information options.
4. Click Apply

Figure 232: Configuring Global Settings for DHCPv6 Snooping

Security > DHCP Snooping6

Step: 1. Configure Global

DHCPv6 Snooping Status ☐ Enabled

DHCPv6 Snooping Option Remote ID ☐ Enabled

DHCPv6 Snooping Option Policy Drop

Apply Revert

DHCPv6 Snooping VLAN Configuration Use the Security > DHCP Snooping6 (Configure VLAN) page to enable or disable DHCPv6 snooping on specific VLANs.

Command Usage

- ◆ When DHCPv6 snooping enabled globally and enabled on a VLAN, DHCPv6 packet filtering will be performed on any untrusted ports within the VLAN.
- ◆ When the DHCPv6 snooping is globally disabled, DHCPv6 snooping can still be configured for specific VLANs, but the changes will not take effect until DHCPv6 snooping is globally re-enabled.
- ◆ When DHCPv6 snooping is enabled globally, and 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)

Web Interface

To configure global settings for DHCPv6 Snooping:

1. Click Security, DHCP Snooping6.
2. Select Configure VLAN from the Step list.
3. Select Add from the Action list.
4. Select a VLAN on which to enable DHCPv6 Snooping.
5. Click Apply

Figure 233: Configuring DHCPv6 Snooping on a VLAN

Security > DHCP Snooping6

Step: 2. Configure VLAN Action: Add

VLAN ID (1-4094) 1

Apply Revert

To show the VLANs for which DHCPv6 Snooping is enabled:

1. Click Security, DHCP Snooping6.
2. Select Configure VLAN from the Step list.
3. Select Show from the Action list.

Figure 234: Showing VLANs Enabled for DHCPv6 Snooping

Security > DHCP Snooping6

Step: 2. Configure VLAN Action: Show

DHCPv6 Snooping VLAN List: total 1

VLAN
1

Apply Revert

Configuring Interfaces for DHCPv6 Snooping

Use the Security > DHCP Snooping6 (Configure Interface) page to configure switch interfaces as trusted or untrusted, and set the maximum number of entries which can be stored in the binding database for an interface.

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.

- ◆ Set all interfaces connected to DHCPv6 servers within the local network or fire wall to trusted, and all other interfaces outside the local network or fire wall to untrusted.
- ◆ When DHCPv6 snooping is enabled globally and enabled on a VLAN, DHCPv6 packet filtering will be performed on any untrusted ports within the VLAN according to the default status, or as specifically configured for an interface.
- ◆ When an untrusted port is changed to a trusted port, all the dynamic DHCPv6 snooping bindings associated with this port are removed.
- ◆ *Additional considerations when the switch itself is a DHCPv6 client* – The port(s) through which it submits a client request to the DHCPv6 server must be configured as trusted.

Parameters

These parameters are displayed:

- ◆ **Interface** – Port or trunk identifier.
- ◆ **Trust Status** – Enables or disables an interface as trusted. (Default: Disabled)
- ◆ **Max Binding** – Sets the maximum number of entries which can be stored in the binding database for an interface. (Range: 1-5; Default: 5)
- ◆ **Current Binding** – Shows the maximum number of entries which can be stored in the binding database for an interface.

Web Interface

To configure the trust status and maximum bindings for DHCPv6 Snooping:

1. Click Security, DHCP Snooping6.
2. Select Configure Interface from the Step list.
3. Set any interfaces within the local network or firewall to trusted.
4. Set the maximum number of entries which can stored for an interface.
5. Click Apply

Figure 235: Configuring the Trust State for DHCPv6 Snooping

Security > DHCP Snooping6

Step: 3. Configure Interface

DHCPv6 Snooping Interface List: Total 28

Interface	Trust Status	Max Binding (1 - 5)	Current Binding
Eth 1/1	☐ Enabled	5	0
Eth 1/2	☐ Enabled	5	0
Eth 1/3	☐ Enabled	5	0
Eth 1/4	☐ Enabled	5	0
Eth 1/5	☐ Enabled	5	0

Displaying DHCPv6 Snooping Binding Information

Use the Security > DHCP Snooping6 (Show Information – Binding) page to display entries in the binding table.

Parameters

These parameters are displayed:

- ◆ **Link-layer Address** – IPv6 link-layer address associated with the entry.
- ◆ **IPv6 Address** – IPv6 address corresponding to the client.
- ◆ **Lifetime** – The time (number of seconds) for which this IPv6 address is leased to the client.
- ◆ **VLAN** – VLAN to which this entry is bound.
- ◆ **Interface** – Port or trunk to which this entry is bound.
- ◆ **Type** – Entry types include:
 - **NA** – Non-temporary address.
 - **TA** – Temporary address.
- ◆ **Clear** – Removes all dynamically learned snooping entries from RAM.

Web Interface

To display the binding table for DHCPv6 Snooping:

1. Click Security, DHCP Snooping6.
2. Select Show Information from the Step list.
3. Click Binding.
4. Use the Clear function if required.

Figure 236: Displaying the Binding Table for DHCPv6 Snooping



	Link-Local Address	IPv6 Address	Lifetime	VLAN	Interface	Type
☐	00-13-4F-AA-18-26	2001:021:143E:5612:a53:c5782:a452:9712	2561998	1	Eth 1/5	NA
☐	8B-12-CF-81-3D-03	2001:0300:1	2561912	1	Eth 1/3	NA
☐	8B-12-CF-81-3D-04	2001:0300:3	2561555	1	Trunk 2	TA

Displaying DHCPv6 Snooping Statistics Use the Security > DHCP Snooping6 (Show Information – Statistics) page to display information on client, server, and relay packets.

Parameters

These parameters are displayed:

- ◆ **State** – Packet states include received, sent and dropped.
- ◆ **Types**
 - **Client Packet** – Includes Solicit, Request, Confirm, Renew, Rebind, Decline, Release and Information-request.
 - **Server Packet** – Includes Advertise, Reply, and Reconfigure.
 - **Relay Packet** – Includes Relay-forward and Relay-reply.

Web Interface

To display statistics for DHCPv6 Snooping:

1. Click Security, DHCP Snooping6.
2. Select Show Information from the Step list.
3. Click Statistics.
4. Use the Clear function if required.

Figure 237: Displaying Statistics for DHCPv6 Snooping



ND Snooping

Neighbor Discovery (ND) Snooping maintains an IPv6 prefix table and user address binding table. These tables can be used for stateless address auto-configuration or for address filtering by IPv6 Source Guard.

ND snooping maintains a binding table in the process of neighbor discovery. When it receives an Neighbor Solicitation (NS) packet from a host, it creates a new binding. If it subsequently receives a Neighbor Advertisement (NA) packet, this means that the address is already being used by another host, and the binding is therefore deleted. If it does not receive an NA packet after a timeout period, the binding will be bound to the original host. ND snooping can also maintain a prefix table used for stateless address auto-configuration by monitoring Router Advertisement (RA) packets sent from neighboring routers.

ND snooping can also detect if an IPv6 address binding is no longer valid. When a binding has been timed out, it checks to see if the host still exists by sending an NS packet to the target host. If it receives an NA packet in response, it knows that the target still exists and updates the lifetime of the binding; otherwise, it deletes the binding.

Usage Guidelines

- ◆ ND snooping must be enabled globally on the switch and on a specific VLAN or a range of VLANs.
- ◆ Once ND snooping is enabled both globally and on the required VLANs, the switch will start monitoring RA messages to build an address prefix table as described below:

- If an RA message is received on an untrusted interface, it is dropped. If received on a trusted interface, the switch adds an entry in the prefix table according to the Prefix Information option in the RA message. The prefix table records prefix, prefix length, valid lifetime, as well as the VLAN and port interface which received the message.
- If an RA message is not received updating a table entry with the same prefix for a specified timeout period, the entry is deleted.
- ◆ Once ND snooping is enabled both globally and on the required VLANs, the switch will start monitoring NS messages to build a dynamic user binding table for use in Duplicate Address Detection (DAD) or for use by other security filtering protocols (e.g., IPv6 Source Guard) as described below:
 - If an NS message is received on an trusted interface, it is forwarded without further processing.
 - If an NS message is received on an untrusted interface, and the address prefix does not match any entry in the prefix table, it drops the packet.
 - If the message does match an entry in the prefix table, it adds an entry to the dynamic user binding table after a fixed delay, and forwards the packet. Each entry in the dynamic binding table includes the link-layer address, IPv6 address, lifetime, as well as the VLAN and port interface which received the message.
 - If an RA message is received in response to the original NS message (indicating a duplicate address) before the dynamic binding timeout period expires, the entry is deleted. Otherwise, when the timeout expires, the entry is dropped if the auto-detection process is not enabled.
 - If the auto-detection process is enabled, the switch periodically sends an NS message to determine if the client still exists. If it does not receive an RA message in response after the configured timeout, the entry is dropped. If the switch receives an RA message before the timeout expires, it resets the lifetime for the dynamic binding, and the auto-detection process resumes.

ND Snooping Global Configuration Use the Security > ND Snooping > Configure Global page to enable ND snooping globally on the switch.

Parameters

These parameters are displayed:

- ◆ **ND Snooping Status** – Enables ND Snooping globally on the switch. (Default: Disabled)
- ◆ **ND Snooping Auto-detect** – Enables automatic validation of dynamic user binding table entries by periodically sending NS messages and awaiting NA replies.

If auto-detection is enabled, the switch periodically sends an NS message to determine if a client listed in the dynamic binding table still exists. If it does not

receive an RA message in response after the configured timeout, the entry is dropped. If the switch receives an RA message before the timeout expires, it resets the lifetime for the dynamic binding, and the auto-detection process resumes. (Default: Disabled)

- ◆ **ND Snooping Retransmit Count** – Sets the number of times the auto-detection process sends an NS message to determine if a dynamic user binding is still valid. (Range: 1-5 seconds; Default: 3 seconds)
- ◆ **ND Snooping Retransmit Interval** – Sets the interval between which the auto-detection process sends NS messages to determine if a dynamic user binding is still valid.

The timeout after which the switch will delete a dynamic user binding if no RA message is received is set to the retransmit count multiplied by the retransmit interval. (Range: 1-10 seconds; Default: 1 second)

- ◆ **ND Snooping Prefix Timeout** – Sets the time to wait for an RA message before deleting an entry in the prefix table. If ND snooping is enabled and an RA message is received on a trusted interface, the switch will add an entry in the prefix table based upon the Prefix Information contained in the message. If an RA message is not received for a table entry with the same prefix for the specified timeout period, the entry is deleted. (Range: 3-1800 seconds; Default: *none set*)

Web Interface

To configure ND Snooping globally for the switch:

1. Click Security, ND Snooping, Configure Global.
2. Enable ND Snooping Status and set other parameters as required.
3. Click Apply.

Figure 238: ND Snooping Global Configuration

Security > ND Snooping

Step: 1. Configure Global

ND Snooping Status ☐ Enabled

ND Snooping Auto-detect ☐ Enabled

ND Snooping Retransmit Count(1-5)

ND Snooping Retransmit Interval(1-10) sec

☐ ND Snooping Prefix Timeout(3-1800)

Apply Revert

ND Snooping VLAN Configuration Use the Security > ND Snooping > Configure VLAN (Add) page to enable ND Snooping on a specified VLAN or range of VLANs.

Parameters

These parameters are displayed:


- ◆ **VLAN ID** - A specific VLAN ID or a consecutive range of VLANs. (Range: 1-4094)

Web Interface

To configure ND Snooping on a VLAN:

1. Click Security, ND Snooping, Configure VLAN.
2. Specify a VLAN ID or range of VLAN IDs.
3. Click Apply.

Figure 239: ND Snooping VLAN Configuration



The screenshot shows a web interface for configuring ND Snooping. The breadcrumb path is "Security > ND Snooping". Below this, there are two dropdown menus: "Step:" set to "2. Configure VLAN" and "Action:" set to "Add". The main configuration area is labeled "VLAN ID (1-4094)" and contains two input boxes separated by a hyphen, representing a range of VLANs. At the bottom right, there are two buttons: "Apply" and "Revert".

Configuring Ports for ND Snooping Use the Security > ND Snooping > Configure Interface page to configure ports as trusted interfaces from which prefix information in RA messages can be added to the prefix table, or NS messages can be forwarded without validation.

Usage Guidelines

- ◆ In general, interfaces facing toward to the network core, or toward routers supporting the Network Discovery protocol, are configured as trusted interfaces.
- ◆ RA messages received from a trusted interface are added to the prefix table and forwarded toward their destination.
- ◆ NS messages received from a trusted interface are forwarded toward their destination. Nothing is added to the dynamic user binding table.

Parameters

These parameters are displayed:

- ◆ **Trust Status** – Enables or disables a port as trusted. (Default: Disabled)
- ◆ **Max Binding** – The maximum number of address entries in the dynamic user binding table which can be bound to a port. (Range: 1-5; Default: 5)

Web Interface

To configure ND Snooping on a port interface:

1. Click Security, ND Snooping, Configure Interface.
2. Set the required ports' Trust Status to enabled.
3. Click Apply.

Figure 240: ND Snooping Interface Configuration

Interface	Trust Status	Max Binding (1 - 5)
Eth 1/1	☐ Enabled	5
Eth 1/2	☐ Enabled	5
Eth 1/3	☐ Enabled	5
Eth 1/4	☐ Enabled	5
Eth 1/5	☐ Enabled	5

Displaying ND Snooping Binding Information

Use the Security > ND Snooping > Show Information (Binding) page to display the ND Snooping binding list.

Parameters

These parameters are displayed:

- ◆ **Mac Address** – The MAC address associated with the entry.
- ◆ **IPv6 Address** – The IPv6 address corresponding to the client.
- ◆ **Lifetime** – The amount of time that the specified IPv6 prefix is advertised as being preferred. The preferred lifetime is counted down in real time.
- ◆ **VLAN** – The VLAN to which this entry is bound.
- ◆ **Interface** – The port or trunk interface to which this entry is bound.

Web Interface

To display the ND Snooping binding list:

1. Click Security, ND Snooping.
2. Select Show Information from the Step list.
3. Click Binding.

Figure 241: Displaying the Binding List for ND Snooping



Displaying ND Snooping Prefix Information

Use the Security > ND Snooping > Show Information (Prefix) page to display the ND Snooping prefix list.

Parameters

These parameters are displayed:

- ◆ **Prefix** – The IPv6 network prefix associated with the entry.
- ◆ **Len** – (Prefix length) A decimal value that indicates how many contiguous bits (from the left) of the IPv6 address comprise the network portion of the address.
- ◆ **Valid-Time** – The amount of time that the specified IPv6 prefix is advertised as being valid.
- ◆ **Expire** – The time before this prefix entry will be removed.
- ◆ **VLAN** – The VLAN associated with the prefix entry.
- ◆ **Interface** – The port or trunk interface associated with the prefix entry.

Web Interface

To display the ND Snooping prefix list:

1. Click Security, ND Snooping.
2. Select Show Information from the Step list.
3. Click Prefix.

Figure 242: Displaying the Prefix List for ND Snooping

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 364](#)). 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 > General 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 364](#)), 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 367](#)), 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:

- ◆ **Port** – The port identifier.
- ◆ **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: 1-10, default: 5; MAC Table: 1-1024, default: 16)

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 364](#)) and static entries set by IP source guard (see [“Configuring Static Bindings for IPv4 Source Guard” on page 389](#)).

Web Interface

To set the IP Source Guard filter for ports:

1. Click Security, IP Source Guard, General.
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 243: Setting the Filter Type for IPv4 Source Guard

Port	Filter Type	Filter Table	ACL Table Max Binding Entry (1-10)	MAC Table Max Binding Entry (1-1024)
1	Disabled	ACL	5	16
2	Disabled	ACL	5	16
3	Disabled	ACL	5	16
4	Disabled	ACL	5	16
5	Disabled	ACL	5	16
6	Disabled	ACL	5	16
7	Disabled	ACL	5	16
8	Disabled	ACL	5	16
9	Disabled	ACL	5	16
10	Disabled	ACL	5	16

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: 1-28/54)

Show

- ◆ **MAC Address** – Physical address associated with the entry.
- ◆ **IP Address** – IP address corresponding to the client.
- ◆ **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 244: 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: 08-10-15-14-00-01 (xx-xx-xx-xx-xx-xx of xxxxxxxxxxxx)

IP Address: 10.2.44.95

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 245: Displaying Static Bindings for IPv4 Source Guard



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.
- ◆ **Type** – Dynamic DHCPv4 binding, stateful address.

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 246: Showing the IPv4 Source Guard Binding Table

Security > IP Source Guard > Dynamic Binding

Query by:

☐ Port

☐ VLAN

☐ MAC Address

☐ IP Address

Dynamic Binding List Total: 3

VLAN	MAC Address	Interface	IP Address	Type
1	00-10-85-F4-00-E1	Unit 1 / Port 2	192.44.96	DHCP
1	00-10-85-F4-00-E2	Unit 1 / Port 4	192.44.97	DHCP
2	00-10-85-F4-00-E3	Unit 1 / Port 7	192.44.98	DHCP

IPv6 Source Guard

IPv6 Source Guard is a security feature that filters IPv6 traffic on non-routed, Layer 2 network interfaces based on manually configured entries in the IPv6 Source Guard table, or dynamic entries in the Neighbor Discovery Snooping table or DHCPv6 Snooping table when either snooping protocol is enabled (refer to the DHCPv6 Snooping commands in the *CLI Reference Guide*). IPv6 source guard can be used to prevent traffic attacks caused when a host tries to use the IPv6 address of a neighbor to access the network. This section describes how to configure IPv6 Source Guard.

Configuring Ports for IPv6 Source Guard

Use the Security > IPv6 Source Guard > Port Configuration page to filter inbound traffic based on the source IPv6 address stored in the binding table.

IPv6 Source Guard is used to filter traffic on an insecure port which receives messages from outside the network or fire wall, and therefore may be subject to traffic attacks caused by a host trying to use the IPv6 address of a neighbor.

Command Usage

- ◆ Setting source guard mode to SIP (Source IP) enables this function on the selected port. Use the SIP option to check the VLAN ID, IPv6 global unicast source IP address, and port number against all entries in the binding table.
- ◆ After IPv6 source guard is enabled on an interface, the switch initially blocks all IPv6 traffic received on that interface, except for ND packets allowed by ND snooping and DHCPv6 packets allowed by DHCPv6 snooping. A port access control list (ACL) is applied to the interface. Traffic is then filtered based upon dynamic entries learned via ND snooping or DHCPv6 snooping, or static addresses configured in the source guard binding table. The port allows only IPv6 traffic with a matching entry in the binding table and denies all other IPv6 traffic.

- ◆ Table entries include a MAC address, IPv6 global unicast address, entry type (Static-IPv6-SG-Binding, Dynamic-ND-Binding, Dynamic-DHCPv6-Binding), VLAN identifier, and port identifier.
- ◆ Static addresses entered in the source guard binding table (using the Static Binding page) are automatically configured with an infinite lease time. Dynamic entries learned via DHCPv6 snooping are configured by the DHCPv6 server itself.
- ◆ If IPv6 source guard is enabled, an inbound packet's source IPv6 address will be checked against the binding table. If no matching entry is found, the packet will be dropped.
- ◆ Filtering rules are implemented as follows:
 - If ND snooping and DHCPv6 snooping are disabled, IPv6 source guard will check the VLAN ID, source IPv6 address, and port number. If a matching entry is found in the binding table and the entry type is static IPv6 source guard binding, the packet will be forwarded.
 - If ND snooping or DHCP snooping is enabled, IPv6 source guard will check the VLAN ID, source IP address, and port number. If a matching entry is found in the binding table and the entry type is static IPv6 source guard binding, dynamic ND snooping binding, or dynamic DHCPv6 snooping binding, the packet will be forwarded.
 - If IPv6 source guard is enabled on an interface for which IPv6 source bindings (dynamically learned via ND snooping or DHCPv6 snooping, or manually configured) are not yet configured, the switch will drop all IPv6 traffic on that port, except for ND packets and DHCPv6 packets allowed by DHCPv6 snooping.
 - Only IPv6 global unicast addresses are accepted for static bindings.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier.
- ◆ **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 395).

- 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 247: Setting the Filter Type for IPv6 Source Guard

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 and 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 Binding.
2. Select Add from the Action list.
3. Enter the required bindings for each port.
4. Click Apply

Figure 248: Configuring Static Bindings for IPv6 Source Guard

Security > IPv6 Source Guard > Static Binding

Action: Add ▼

Port: 1 ▼

VLAN: 1 ▼

MAC Address: 00-10-b5-14-00-01 (00-XX-XX-XX-XX-XX or XXXXXXXXXXXX)

IPv6 Address: 2001::1

Apply Revert

To display static bindings for IPv6 Source Guard:

1. Click Security, IPv6 Source Guard, Static Binding.
2. Select Show from the Action list.

Figure 249: 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-b5-14-00-01	Eth 1/0	2001:DB8:2222:7273::26	STA
☐	1	00-10-b5-14-00-02	Eth 1/8	2001:DB8:2222:7273::58	DHCP
☐	2	00-10-b5-14-00-03	Eth 1/7	2001:DB8:2222:7273::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.
 - **STA** – Static IPv6 Source Guard binding.

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 250: Showing the IPv6 Source Guard Binding Table

Security > IPv6 Source Guard > Dynamic Binding

Query by:

☐ Port

☐ VLAN

☐ MAC Address (00-10-85-F4-00-01 OF 000000000000)

☐ IPv6 Address

Dynamic Binding List Total: 5

VLAN	MAC Address	Interface	IPv6 Address	Type
1	00-10-85-F4-00-01	Eth 1/0	2001::de	ND
1	00-10-85-F4-00-02	Eth 1/4	2001::de	ND
2	00-10-85-F4-00-03	Pcs 1	2001::de	ND

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 367](#)). 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 341](#)).

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: 20)
- ◆ **Log Interval** – The interval at which log messages are sent. (Range: 0-86400 seconds; Default: 10 seconds)

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 251: 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 'ARP Inspection Status' is checked and labeled 'Enabled'. Under 'ARP Inspection Validation', four checkboxes are present: 'Dst-MAC', 'IP', 'Allow Zeroes', and 'Src-MAC'. The 'Log Message Number (0-256)' is set to 20, and the 'Log Interval (0-86400)' is set to 10 seconds. 'Apply' and 'Revert' buttons are located at the bottom right of the configuration area.

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 341](#)).
- ◆ 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 252: Configuring VLAN Settings for ARP Inspection

Security > ARP Inspection

Step: 2: Configure VLAN

ARP Inspection VLAN List: Total: 3

VLAN	DAI Status	ACL Name	ACL Status
1	☒ Enabled	☒ arp_acl_1	☒ Static
2	☒ Enabled	☒ arp_acl_1	☐ Static
3	☒ Enabled	☒ arp_acl_2	☐ 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: 0-750 pps; Default: 15 pps)
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 253: Configuring Interface Settings for ARP Inspection

Port	Trust Status	Packet Rate Limit (0-750 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 23: ARP Inspection Statistics


Parameter	Description
Received ARP packets before ARP inspection rate limit	Count of ARP packets received but not exceeding the ARP Inspection rate limit.
Dropped ARP packets in the process of ARP inspection rate limit	Count of ARP packets exceeding (and dropped by) ARP rate limiting.
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 254: Displaying Statistics for ARP Inspection



The screenshot shows the 'Security > ARP Inspection' page. At the top, there are two dropdown menus: 'Step: 4. Show Information' and 'Action: Show Statistics'. Below these, a table displays various statistics for ARP inspection.

Received ARP packets before ARP inspection rate limit	1000
Dropped ARP packets in processing ARP inspection rate limit	5
Total ARP packets processed by ARP inspection	200
ARP packets dropped by additional validation (Src-MAC)	300
ARP packets dropped by additional validation (Dst-MAC)	2000
ARP packets dropped by additional validation (IP)	100
ARP packets dropped by ARP ACLs	5
ARP packets dropped by DHCP snooping	5

Displaying the ARP Inspection Log

Use the Security > ARP Inspection (Show Information - Show Log) page to show information about entries stored in the log, including the associated VLAN, port, and address components.

Parameters

These parameters are displayed:

Table 24: ARP Inspection Log

Parameter	Description
VLAN ID	The VLAN where this packet was seen.
Port	The port where this packet was seen.
Src. IP Address	The source IP address in the packet.
Dst. IP Address	The destination IP address in the packet.
Src. MAC Address	The source MAC address in the packet.
Dst. MAC Address	The destination MAC address in the packet.

Web Interface

To display the ARP Inspection log:

1. Click Security, ARP Inspection.
2. Select Show Information from the Step list.
3. Select Show Log from the Action list.

Figure 255: Displaying the ARP Inspection Log

VLAN ID	Port	Src. IP Address	Dest. IP Address	Src. MAC Address	Dest. 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-33-75-55-8B	AD-BE-CE-DD-4E-1F

Application Filter

Use the Security > Application Filter page to forward CDP or PVST packets.

Command Usage

If this feature is not enabled, the switch will handle CDP or PVST packets as normal packets. In other words, they are forwarded to other ports in the same VLAN that are also configured to forward the specified packet type.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier (Range: 1-28/54)
- ◆ **CDP** – Cisco Discovery Protocol
- ◆ **PVST** – Per-VLAN Spanning Tree

Web Interface

To discard or forward CDP or PVST packets:

1. Click Security, Application Filter.
2. Set the packet type to be discarded or forwarded as required.
3. Click Apply.

Figure 256: Configuring Discarding or Forwarding of CDP/PVST Packets

Port	CDP	PVST
1	Default	Default
2	Default	Default
3	Default	Default
4	Default	Default
5	Default	Default

Basic Administration Protocols

This chapter describes basic administration tasks including:

- ◆ **Event Logging** – Sets conditions for logging event messages to system memory or flash memory, configures conditions for sending trap messages to remote log servers, and configures trap reporting to remote hosts using Simple Mail Transfer Protocol (SMTP).
- ◆ **Link Layer Discovery Protocol (LLDP)** – Configures advertisement of basic information about the local switch, or discovery of information about neighboring devices on the local broadcast domain.
- ◆ **Power over Ethernet** – Sets the priority and power budget for each port.
- ◆ **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.
- ◆ **Time Range** – Sets a time range during which various functions are applied, including applied ACLs.
- ◆ **Ethernet Ring Protection Switching (ERPS)** – Configures a protection switching mechanism and protocol for Ethernet layer network rings.
- ◆ **MLAG** - Configures multi-chassis link aggregation group (MLAG).
- ◆ **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.

- ◆ **Loopback Detection (LBD)** – Detects general loopback conditions caused by hardware problems or faulty protocol settings.
- ◆ **Smart Pair Configuration** – 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 25: Logging Levels

Level	Severity Name	Description
7	Debugging	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 257: Configuring Settings for System Memory Logs

Administration > Log > System

Step: 2. Configure Global

Status: ☒ Enabled

History Flash Level: 3 - Error

History 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 258: Showing Error Messages Logged to System Memory



Remote Log Configuration

Use the Administration > Log > Remote page to send log messages to syslog servers or other management stations. You can also limit the event messages sent to only those messages below a specified level.

Parameters

These parameters are displayed:

- ◆ **Remote Log Status** – Enables/disables the logging of debug or error messages to the remote logging process. (Default: Disabled)
- ◆ **Logging Facility** – Sets the facility type for remote logging of syslog messages. There are eight facility types specified by values of 16 to 23. The facility type is used by the syslog server to dispatch log messages to an appropriate service.

The attribute specifies the facility type tag sent in syslog messages (see RFC 3164). This type has no effect on the kind of messages reported by the switch. However, it may be used by the syslog server to process messages, such as sorting or storing messages in the corresponding database. (Range: 16-23, Default: 23)
- ◆ **Logging Trap Level** – Limits log messages that are sent to the remote syslog server for all levels up to the specified level. For example, if level 3 is specified, all messages from level 0 to level 3 will be sent to the remote server. (Range: 0-7, Default: 7)
- ◆ **Server IP Address** – Specifies the IPv4 or IPv6 address of a remote server which will be sent syslog messages.

- ◆ **Port** - Specifies the UDP port number used by the remote server.
(Range: 1-65535; Default: 514)

Web Interface

To configure the logging of error messages to remote servers:

1. Click Administration, Log, Remote.
2. Enable remote logging, specify the facility type to use for the syslog messages, and enter the IP address of the remote servers.
3. Click Apply.

Figure 259: Configuring Settings for Remote Logging of Error Messages

Administration > Log > Remote

Remote Log Status: ☒ Enabled

Logging Facility: 23 - Local use 7

Logging Trap Level: 0 - System unusable

Server IP Address 1: 192.168.0.4 Port: 514

Server IP Address 2: Port:

Server IP Address 3: Port:

Server IP Address 4: Port:

Server IP Address 5: Port:

Apply Revert

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 411](#)) 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 705](#)), and one or more DNS servers specified (see [“Configuring a List of Name Servers” on page 708](#), or [“Configuring Static DNS Host to Address Entries” on page 709](#)).

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 260: Configuring SMTP Alert Messages

Administration > Log > SMTP

SMTP Status	☒ Enabled
Severity	3 - Error
E-mail Source Address	big-wheels@matel.com
E-mail Destination Address 1	chris@matel.com
E-mail Destination Address 2	
E-mail Destination Address 3	
E-mail Destination Address 4	
E-mail Destination Address 5	
Server IP Address 1	192.168.1.4
Server IP Address 2	
Server IP Address 3	

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 261: Configuring LLDP Timing Attributes

The screenshot shows the 'Administration > LLDP' configuration page. At the top, there's a 'Steps' dropdown menu with '1. Configure Global' selected. Below this, the 'LLDP' section has a checkbox labeled 'Enabled' which is checked. Underneath, there are six configuration items, each with a text input field and a unit 'sec':
- 'Transmission Interval (5-32768)' with a value of '30'.
- 'Hold Time Multiplier (2-10)' with a value of '4'.
- 'Delay Interval (1-8192)' with a value of '2'.
- 'Reinitialization Delay (1-60)' with a value of '2'.
- 'Notification Interval (5-3600)' with a value of '5'.
- 'MED Fast Start Count (1-10)' with a value of '4'.
At the bottom, there is a note: 'Note: The Transmission Interval must be greater than or equal to 4 times the Delay Interval.' and two buttons: 'Apply' and 'Revert'.

Configuring LLDP Interface Attributes

Use the Administration > LLDP (Configure Interface - Configure General) page to specify the message attributes for individual interfaces, including whether messages are transmitted, received, or both transmitted and received, whether SNMP notifications are sent, and the type of information advertised.

Parameters

These parameters are displayed:

- ◆ **Admin Status** – Enables LLDP message transmit and receive modes for LLDP Protocol Data Units. (Options: Tx only, Rx only, TxRx, Disabled; Default: TxRx)
- ◆ **SNMP Notification** – Enables the transmission of SNMP trap notifications about LLDP and LLDP-MED changes. (Default: Disabled)

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 461](#).

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 IP 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.

- **Management IPv6 Address** – The management address protocol packet includes the IPv6 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: Disabled)
- **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 74](#). (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 187](#)). (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 161](#)). (Default: Enabled)
 - **VLAN Name** – The name of all VLANs to which this interface has been assigned (see [“IEEE 802.1Q VLANs” on page 161](#)). (Default: Enabled)
 - **Port and Protocol VLAN ID** – The port-based protocol VLANs configured on this interface (see [“Protocol VLANs” on page 187](#)). (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 76](#) 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)
 - **Extended Power** – (ECS4150-28P) This option advertises extended Power-over-Ethernet capability details, such as power availability from the switch, and power state of the switch, including whether the switch is operating from primary or backup power (the Endpoint Device could use this information to decide to enter power conservation mode). (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 the DHCP server.
 - Location of the network element believed to be the closest to client.
 - Location of the 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 262: Configuring LLDP Interface Attributes

The screenshot shows the 'Administration > LLDP' configuration page. The 'Interface' dropdown is set to 'Port'. The 'Admin Status' is 'Tx/Rx'. The 'SNMP Notification' and 'MED Notification' are both 'Enabled'. Under 'Basic Optional TLVs', 'Management IP Address', 'System Description', 'Port Describes', and 'System Capabilities' are checked. Under '802.1 Organizationally Specific TLVs', 'Protocol Identity', 'VLAN ID', 'VLAN Name', and 'Port and Protocol VLAN ID' are checked. Under '802.3 Organizationally Specific TLVs', 'Link Aggregation', 'Max Frame Size', and 'MAC/PHY Configuration/Status' are checked. Under 'MED TLV's', 'Capabilities', 'Network Policy', 'Inventory', and 'Location' are checked. The 'MED-Location Civic Address' section shows 'Country' as 'TW' 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 26: 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

Table 26: LLDP MED Location CA Types (Continued)

CA Type	Description	CA Value Example
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 or Modify 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 263: Configuring the Civic Address for an LLDP Interface

The screenshot shows a web interface titled "Administration > LLDP". Below the title, there is a "Step:" dropdown menu set to "2. Configure Interface" and an "Action:" dropdown menu set to "Add CA-Type". The main configuration area has three sections: "Interface" with radio buttons for "Port" (selected) and "Trunk"; "CA-Type (0-255)" with a text input field containing "1"; and "CA-Value" with a text input field containing "California". At the bottom right, there are "Apply" and "Revert" buttons.

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 264: Showing the Civic Address for an LLDP Interface



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 27: 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 74](#)).

- ◆ **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 28: System Capabilities

ID Basis	Reference
Other	—
Repeater	IETF RFC 2108
Bridge	IETF RFC 2674
WLAN Access Point	IEEE 802.11 MIB
Router	IETF RFC 1812
Telephone	IETF RFC 2011
DOCSIS cable device	IETF RFC 2669 and IETF RFC 2670
End Station Only	IETF RFC 2011

- ◆ **System Capabilities Enabled** – The primary function(s) of the system which are currently enabled. Refer to the preceding table.
- ◆ **Management Address** – The management address associated with the local system. If no management address is available, the address should be the MAC address for the CPU or for the port sending this advertisement.

Interface Settings

The attributes listed below apply to both port and trunk interface types. When a trunk is listed, the descriptions apply to the first port of the trunk.

- ◆ **Port/Trunk Description** – A string that indicates the port or trunk description. If RFC 2863 is implemented, the ifDescr object should be used for this field.
- ◆ **Port/Trunk ID** – A string that contains the specific identifier for the port or trunk from which this LLDPDU was transmitted.

Interface Details

The attributes listed below apply to both port and trunk interface types. When a trunk is listed, the descriptions apply to the first port of the trunk.

- ◆ **Local Port/Trunk** – Local interface on this switch.

- ◆ **Port/Trunk ID Type** – There are several ways in which a port may be identified. A port ID subtype is used to indicate how the port is being referenced in the Port ID TLV.

Table 29: Port ID Subtype

ID Basis	Reference
Interface alias	IfAlias (IETF RFC 2863)
Chassis component	EntPhysicalAlias when entPhysClass has a value of 'chassis(3)' (IETF RFC 2737)
Port component	EntPhysicalAlias when entPhysicalClass has a value 'port(10)' or 'backplane(4)' (IETF RFC 2737)
MAC address	MAC address (IEEE Std 802-2001)
Network address	networkAddress
Interface name	ifName (IETF RFC 2863)
Agent circuit ID	agent circuit ID (IETF RFC 3046)
Locally assigned	locally assigned

- ◆ **Port/Trunk ID** – A string that contains the specific identifier for the local interface based on interface subtype used by this switch.
- ◆ **Port/Trunk Description** – A string that indicates the port or trunk description. If RFC 2863 is implemented, the ifDescr object should be used for this field.
- ◆ **MED Capability** – The supported set of capabilities that define the primary function(s) of the interface:
 - LLDP-MED Capabilities
 - Network Policy
 - Location Identification
 - 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 265: 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	ISAC Address
Chassis ID	E4-9D-73-1B-5E-00
System Name	
System Description	ECSE150-28P
System Capabilities Supported	Bridge, Router
System Capabilities Enabled	Bridge, Router
Management Address	192.168.1.1 (IPv4)

Figure 266: 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	PortID
1	Ethernet Port on unit 1, port 1	88-3C-83-8C-44-B3
2	Ethernet Port on unit 1, port 2	88-3C-83-8C-44-B4
3	Ethernet Port on unit 1, port 3	88-3C-83-8C-44-B5
4	Ethernet Port on unit 1, port 4	88-3C-83-8C-44-B6
5	Ethernet Port on unit 1, port 5	88-3C-83-8C-44-B7

Figure 267: 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

Port ID Type	ISAC Address
Port ID	E4-9D-73-1B-5E-B1
Port Description	Ethernet Port on unit 1, port 1
MED Capability	LLDP-MED Capabilities, Network Policy, Location Identification, Extended Power via MDI - PSE, 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 27, "Chassis ID Subtype," on page 424.](#))
- ◆ **Chassis ID** – An octet string indicating the specific identifier for the particular chassis in this system.
- ◆ **Port Type** – Indicates the basis for the identifier that is listed in the Port ID field. See [Table 29, "Port ID Subtype," on page 426.](#)
- ◆ **Port ID** – A string that contains the specific identifier for the port from which this LLDPDU was transmitted.
- ◆ **Port Description** – A string that indicates the port's description. If RFC 2863 is implemented, the ifDescr object should be used for this field.
- ◆ **System Name** – A string that indicates the system's assigned name.
- ◆ **System Description** – A textual description of the network entity.

- ◆ **System Capabilities Supported** – The capabilities that define the primary function(s) of the system. (See [Table 28, "System Capabilities," on page 425.](#))
- ◆ **System Capabilities Enabled** – The primary function(s) of the system which are currently enabled. (See [Table 28, "System Capabilities," on page 425.](#))
- ◆ **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 30: 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

Table 30: Remote Port Auto-Negotiation Advertised Capability (Continued)

Bit	Capability
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 Capability ¹⁰

- ◆ **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.

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

Port Details – Network Policy¹⁰

- ◆ **Application Type** – The primary application(s) defined for this network policy:
 - Voice
 - Voice Signaling
 - Guest Signaling
 - Guest Voice Signaling
 - Softphone Voice
 - Video Conferencing
 - Streaming Video
 - Video Signaling
- ◆ **Tagged Flag** – Indicates whether the specified application type is using a tagged or untagged VLAN.
- ◆ **Layer 2 Priority** – The Layer 2 priority to be used for the specified application type. This field may specify one of eight priority levels (0-7), where a value of 0 represents use of the default priority.
- ◆ **Unknown Policy Flag** – Indicates that an endpoint device wants to explicitly advertise that this policy is required by the device, but is currently unknown.
- ◆ **VLAN ID** – The VLAN identifier (VID) for the port as defined in IEEE 802.1Q. A value of zero indicates that the port is using priority tagged frames, meaning that only the IEEE 802.1D priority level is significant and the default PVID of the ingress port is used instead.
- ◆ **DSCP Value** – The DSCP value to be used to provide Diffserv node behavior for the specified application type. This field may contain one of 64 code point values (0-63). A value of 0 represents use of the default DSCP value as defined in RFC 2475.

Port Details – Location Identification¹⁰

- ◆ **Location Data Format** – Any of these location ID data formats:
 - Coordinate-based LCI¹¹ – Defined in RFC 3825, includes latitude resolution, latitude, longitude resolution, longitude, altitude type, altitude resolution, altitude, and datum.
 - Civic Address LCI¹¹ – Includes What, Country code, CA type, CA length and CA value. “What” is described as the field entry “Device entry refers to” under “[Configuring LLDP Interface Attributes](#).” The other items and described under “[Configuring LLDP Interface Civic-Address](#).”

11. Location Configuration Information

- ECS ELIN – Emergency Call Service Emergency Location Identification Number supports traditional PSAP-based Emergency Call Service in North America.
- ◆ **Country Code** – The two-letter ISO 3166 country code in capital ASCII letters. (Example: DK, DE or US)
- ◆ **What** – The type of device to which the location applies as described for the field entry “Device entry refers to” under “[Configuring LLDP Interface Attributes](#).”

Port Details – Extended Power-via-MDI

- ◆ **Power Type** – Power Sourcing Entity (PSE) or Power Device (PD).
- ◆ **Power Priority** – Shows power priority for a port. (Unknown, Low, High, Critical)
- ◆ **Power Source** – Shows information based on the type of device:
 - **PD** – Unknown, PSE, Local, PSE and Local
 - **PSE** – Unknown, Primary Power Source, Backup Power Source - Power conservation mode
- ◆ **Power Value** – The total power in watts required by a PD device from a PSE device, or the total power a PSE device is capable of sourcing over a maximum length cable based on its current configuration. This parameter supports a maximum power required or available value of 102.3 Watts to allow for future expansion. (Range: 0 - 102.3 Watts)

Port Details – Inventory¹⁰

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

The figure consists of two screenshots of a web interface for displaying LLDP information.

The top screenshot shows the "Administration > LLDP" page. The "Step:" dropdown is set to "4. Show Remote Device Information". Below this, there are radio buttons for "Port", "Port Details", "Trunk", and "Trunk Details", with "Port" selected. A table titled "LLDP Remote Device Port List" shows a total of 1 entry. The table has four columns: "Local Port", "Chassis ID", "Port ID", and "System Name". The single entry shows "1" for Local Port, "10-72-0F-32-00-F0" for Chassis ID, "10-72-0F-32-00-F0" for Port ID, and "RDS3" for System Name.

The bottom screenshot shows the same page with the "Port" and "Remote Index" dropdowns set to "1" and "2" respectively. The "Query" button is visible at the bottom right.

Figure 269: Displaying Remote Device Information for LLDP (Port Details)

Administration > LLDP

Step: 4. Show Remote Device Information

☐ Port
 ☒ Port Details
 ☐ Trunk
 ☐ Trunk Details

Port: 1
 Remote Index: 2
 Query

LLDP Remote Device Port Information

Local Port: 1
 Chassis Type: 00-06-03-06-0B-20
 Chassis ID: 1
 Port Type: MAC Address
 Port ID: 00-06-03-06-0B-20
 Port Description: Ethernet Port on unit 1, port 10
 System Name: Ethernet Switch
 System Description: Bridge
 System Capabilities Supported: ECS4150-28T
 System Capabilities Enabled: Bridge

Management Address List: Total: 3

Address	Address Type
192.168.2.3	IPv4 Address
2000::2009	IPv6 Address
08-01-02-03-04-05	MAC 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 Adm-Capability: 0000

Remote Port Auto-Neg Status: Enabled
 Remote Port MAU Type: S

802.3 Extension Power Information

Remote Power Class: PSE
 Remote Power BDM Status: Enabled
 Remote Power Pairs: Spare

Remote Power BDR Supported: Yes
 Remote Power Pair Controlable: No
 Remote Power Classification: Class1

802.3 Extension Trunk Information

Remote Link Aggregation Capable: Yes
 Remote Link Port ID: 0

Remote Link Aggregation Status: Disabled

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 270: 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 Tag	Disabled	VLAN ID	7
Layer 2 Priority	2	DSCP Value	62

Location Identification

Location Data Format	Coordinate-based LCI		
Country Code	TVL	What	2

Location Identification

Location Data Format	Device 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	LN16230092
Manufacturer 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 271: Displaying LLDP Device Statistics (General)

The screenshot shows the 'Administration > LLDP' page. The 'Step' dropdown is set to '5. Show Device Statistics'. Under the 'General' tab, the 'LLDP Device Statistics' section displays the following data:

Statistic	Value
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 272: Displaying LLDP Device Statistics (Port)

The screenshot shows the 'Administration > LLDP' page. The 'Step' dropdown is set to '5. Show Device Statistics'. Under the 'Port' tab, the 'LLDP Device Port Statistics' section displays the following data:

Statistic	Value
Frames Discarded	0
Frames Invalid	0
Frames Received	97
Frames Sent	104
TLVs Unrecognized	0
TLVs Discarded	0
Neighbor Ageouts	0

A 'Refresh' button is located at the bottom right of the statistics section.

Power over Ethernet

The switch can provide DC power to a wide range of connected devices, eliminating the need for an additional power source and cutting down on the amount of cables attached to each device. Once configured to supply power, an automatic detection process is initialized by the switch that is authenticated by a PoE signature from the connected device. Detection and authentication prevent damage to non-compliant devices (prior to IEEE 802.3af).

The switch's power management enables individual port power to be controlled within the switch's power budget. Port power can be automatically turned on and off for connected devices, and a per-port power priority can be set so that the switch never exceeds its power budget. When a device is connected to a switch port, its power requirements are detected by the switch before power is supplied. If the power required by a device exceeds the power budget of the port or the whole switch, port power priority settings are used to control the supplied power.

Ports can be set to one of three power priority levels, critical, high or low. To control the power supply within the switch's budget, ports set at critical to high priority have power enabled in preference to those ports set at low priority. For example, when a device connected to a port is set to critical priority, the switch supplies the required power, if necessary by denying power to ports set for a lower priority.

Setting the Switch's Overall PoE Power Budget

Use the Administration > PoE > PSE (Global Information) page to view the maximum PoE power budget for the switch (power available to all RJ-45 ports). The maximum power budget is fixed at the maximum available setting, which prevents overload conditions at the power source. If the power demand from devices connected to the switch exceeds the power budget, the switch uses port power priority settings to limit the supplied power.

Parameters

These parameters are displayed:

- ◆ **PoE Maximum Available Power** – The power budget for the switch. If devices connected to the switch require more power than the switch budget, the port power priority settings are used to control the supplied power.
(Fixed: 370000 milliwatts for ECS4150-28P, 740000 milliwatts for ECS4150-54P)
- ◆ **PoE Maximum Allocation Power** – The maximum power allocated to switch ports. (Fixed: 370000 milliwatts for ECS4150-28P, 740000 milliwatts for ECS4150-54P)
- ◆ **System Operation Status** – Status of the PoE power service provided to the switch ports.
- ◆ **PoE Power Consumption** – The amount of power being consumed by PoE devices connected to the switch.
- ◆ **Software Version** – The version of software running on the PoE controller subsystem in the switch.

Web Interface

To view the overall PoE power budget for switch:

1. Click Administration, PoE.
2. Select Global Information from the Step list.

Figure 273: Setting the Switch's PoE Budget

The screenshot shows a web-based configuration interface for a switch's PoE settings. The breadcrumb navigation at the top reads 'Administration > PoE > PSE'. Below this, a 'Step:' dropdown menu is set to '1. Global Information'. The main content area displays five configuration items in a table-like format:

PoE Maximum Available Power	370000 milliwatts using main power
PoE Maximum Allocation Power	370000 milliwatts
System Operation Status	Off
PoE Power Consumption	0 milliwatts
Software Version	Version 0007 (Hex)

Setting the Port PoE Power Budget Use the Administration > PoE (Configure Interface) page to set the PoE power allocation priority and maximum power provided to a port.

Command Usage

- ◆ The ECS4150-28P and ECS4150-54P switches support the IEEE 802.3af PoE and IEEE 802.3at-2009 PoE+ standards, and the ECS4150-54P switch also supports the IEEE 802.3bt PoE++ standard on ports 41-48.
- ◆ The total PoE power delivered by all ports cannot exceed the maximum power budget of the switch. The maximum number of ports which can supply power simultaneously at the specified levels are shown in the following table.

Table 31: Maximum Number of Ports Providing Simultaneous Power

Maximum Port Power	ECS4150-28P Max Ports	ECS4150-54P Max Ports
4 W (Class 1)	24	48
7 W (Class 2)	24	48
15.4 W (Class 3)	24	48
30 W (Class 4)	12	24
45 W (Class 5)	-	8 (Ports 41-48 only)
60 W (Class 6)	-	8 (Ports 41-48 only)
75 W (Class 7)	-	8 (Ports 41-48 only)
90 W (Class 8)	-	8 (Ports 41-48 only)

- ◆ If a device is connected to a switch port and the switch detects that it requires more than the power budget set for the port, no power is supplied to the device (i.e., port power remains off).
- ◆ If the power demand from devices connected to all switch ports exceeds the power budget set for the switch, the port power priority settings are used to control the supplied power.
 - If a device connected to a low-priority port causes the switch to exceed its budget, power will be dropped from low-priority ports starting in reverse sequence from the highest RJ-45 port number.

For example, with all ports low priority and ports 2-13 on, when you connect port 1, power to port 13 will turn off and power to port 1 will turn on.

- If a device is connected to a critical or high-priority port that would cause the switch to exceed its power budget as determined during bootup, power is provided to the port only if the switch can drop power to one or more lower-priority ports and thereby remain within its overall budget. Power will be dropped from lower-priority ports starting in reverse sequence from the highest RJ-45 port number.

For example, if ports 2 to 14 request 30W and ports 2-12 are low priority with ports 13-14 high priority, port 12 will turn off and ports 2-11 and 13-14 will turn on.

- If priority is not set for any ports, and PoE consumption exceeds the maximum power provided by the switch, power is shut down starting in reverse sequence from the highest RJ-45 port number.

Parameters

These parameters are displayed:

- ◆ **Port** – The port number on the switch. (Range: 1-24/48)
- ◆ **Admin Status** – Enables PoE power on a port. Power is automatically supplied when a device is detected on a port, providing that the power demanded does not exceed the switch or port power budget. (Default: Enabled)
- ◆ **Mode** – Shows whether or not PoE power is being supplied to a port.
- ◆ **Time Range Name** – Name of a time range. If a time range is set, then PoE will be provided to an interface during the specified period.
- ◆ **Time Range Status** – Indicates if a time range has been applied to an interface, and whether it is currently active or inactive.
- ◆ **Class** – The power class of a connected device.
- ◆ **Priority** – Sets the power priority for a port. (Options: Low, High or Critical; Default: Low)
- ◆ **Power Allocation** – Sets the power budget for a port.
ECS4150-28P: Range: 3000-30000 milliwatts; Default: 30000 milliwatts
ECS4150-54P Ports 1-40: Range: 3000-30000 milliwatts; Default: 30000 milliwatts
ECS4150-54P Ports 41-48: Range: 3000-90000 milliwatts; Default: 90000 milliwatts
- ◆ **Power Consumption** – Current power consumption on a port.

Web Interface

To set the PoE power budget for a port:

1. Click Administration, PoE, PSE.
2. Select Configure Interface from the Step list.
3. Enable PoE power on selected ports. Set the priority and the power budget.
4. Click Apply.

Figure 274: Setting a Port's PoE Budget

Administration > PoE > PSE

Step: 2. Configure Interface

PoE Port List Total: 24

Port	Admin Status	Mode	Time Range Name	Time Range Status	Class	Priority	Power Allocation	Power Consumption (mWatts)
1	☒ Enabled	off		NA	-	Low	30000	0
2	☒ Enabled	off		NA	-	Low	30000	0
3	☒ Enabled	off		NA	-	Low	30000	0
4	☒ Enabled	off		NA	-	Low	30000	0
5	☒ Enabled	off		NA	-	Low	30000	0
6	☒ Enabled	off		NA	-	Low	30000	0
7	☒ Enabled	off		NA	-	Low	30000	0
8	☒ Enabled	off		NA	-	Low	30000	0
9	☒ Enabled	off		NA	-	Low	30000	0
10	☒ Enabled	off		NA	-	Low	30000	0

Apply Reset

Simple Network Management Protocol

Simple Network Management Protocol (SNMP) is a communication protocol designed specifically for managing devices on a network. Equipment commonly managed with SNMP includes switches, routers and host computers. SNMP is typically used to configure these devices for proper operation in a network environment, as well as to monitor them to evaluate performance or detect potential problems.

Managed devices supporting SNMP contain software, which runs locally on the device and is referred to as an agent. A defined set of variables, known as managed objects, is maintained by the SNMP agent and used to manage the device. These objects are defined in a Management Information Base (MIB) that provides a standard presentation of the information controlled by the agent. SNMP defines both the format of the MIB specifications and the protocol used to access this information over the network.

The switch includes an onboard agent that supports SNMP versions 1, 2c, and 3. This agent continuously monitors the status of the switch hardware, as well as the traffic passing through its ports. A network management station can access this information using network management software. Access to the onboard agent from clients using SNMP v1 and v2c is controlled by community strings. To communicate with the switch, the management station must first submit a valid community string for authentication.

Access to the switch from clients using SNMPv3 provides additional security features that cover message integrity, authentication, and encryption; as well as controlling user access to specific areas of the MIB tree.

The SNMPv3 security structure consists of security models, with each model having its own security levels. There are three security models defined, SNMPv1, SNMPv2c, and SNMPv3. Users are assigned to “groups” that are defined by a security model and specified security levels. Each group also has a defined security access to set of MIB objects for reading and writing, which are known as “views.” The switch has a default view (all MIB objects) and default groups defined for security models v1 and v2c. The following table shows the security models and levels available and the system default settings.

Table 32: SNMPv3 Security Models and Levels

Model	Level	Group	Read View	Write View	Notify View	Security
v1	noAuthNoPriv	public (read only)	defaultview	none	none	Community string only
v1	noAuthNoPriv	private (read/write)	defaultview	defaultview	none	Community string only
v1	noAuthNoPriv	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	Community string only
v2c	noAuthNoPriv	public (read only)	defaultview	none	none	Community string only
v2c	noAuthNoPriv	private (read/write)	defaultview	defaultview	none	Community string only
v2c	noAuthNoPriv	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	Community string only
v3	noAuthNoPriv	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	A user name match only
v3	AuthNoPriv	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	Provides user authentication via MD5 or SHA algorithms
v3	AuthPriv	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	<i>user defined</i>	Provides user authentication via MD5 or SHA algorithms and data privacy using DES 56-bit encryption



Note: The predefined default groups and view can be deleted from the system. You can then define customized groups and views for the SNMP clients that require access.

Command Usage

Configuring SNMPv1/2c Management Access

To configure SNMPv1 or v2c management access to the switch, follow these steps:

1. Use the Administration > SNMP (Configure Global) page to enable SNMP on the switch, and to enable trap messages. Refer to: [“Configuring Global Settings for SNMP” on page 445](#)
2. Use the Administration > SNMP (Configure Community - Add) page to configure the community strings authorized for management access. Refer to: [“Setting Community Access Strings” on page 445](#)
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. Refer to: [“Specifying Trap Managers” on page 461](#)

Configuring SNMPv3 Management Access

1. Use the Administration > SNMP (Configure Global) page to enable SNMP on the switch, and to enable trap messages. Refer to: [“Configuring Global Settings for SNMP” on page 445](#)
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. Refer to: [“Specifying Trap Managers” on page 461](#)
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. Refer to: [“Setting the Local Engine ID” on page 447](#)
4. Use the Administration > SNMP (Configure View) page to specify read and write access views for the switch MIB tree. Refer to: [“Setting SNMPv3 Views” on page 449](#)
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). Refer to: [“Configuring SNMPv3 Groups” on page 452](#)
6. Use the Administration > SNMP (Configure Group) page to assign SNMP users to groups, along with their specific authentication and privacy passwords. Refer to: [“Configuring Local SNMPv3 Users” on page 456](#)

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**¹² – Issues a notification message to specified IP trap managers whenever an invalid community string is submitted during the SNMP access authentication process. (Default: Enabled)

Web Interface

To configure global settings for SNMP:

1. Click Administration, SNMP.
2. Select Configure Global from the Step list.
3. Enable SNMP and the required trap types.
4. Click Apply

Figure 275: Configuring Global Settings for SNMP



Setting Community Access Strings

Use the Administration > SNMP (Configure 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)

12. 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 449).

◆ **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 Community from the Step list.
3. Select Add 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 276: Setting Community Access Strings

The screenshot shows the 'Administration > SNMP' web interface. At the top, there's a breadcrumb trail. Below it, a 'Step:' dropdown is set to '2. Configure Community' and an 'Action:' dropdown is set to 'Add'. The main form has two fields: 'Community String' with the value 'spiderman' and 'Access Mode' with a dropdown set to 'Read/Write'. At the bottom right, there are 'Apply' and 'Revert' buttons.

To show the community access strings:

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

Figure 277: Showing Community Access Strings

The screenshot shows the 'Administration > SNMP' web interface with the 'Action:' dropdown set to 'Show'. Below the form, there's a table titled 'SNMP Community String List: Total: 3'. The table has two columns: 'Community String' and 'Access Mode'. There are three rows of data. At the bottom right, there are 'Delete' and 'Revert' buttons.

	Community String	Access Mode
☐	public	Read-Only
☐	private	Read/Write
☐	spiderman	Read/Write

Setting the Local Engine ID Use the Administration > SNMP (Configure Engine - Set Engine ID) page to change the local engine ID. An SNMPv3 engine is an independent SNMP agent that resides on the switch. This engine protects against message replay, delay, and redirection. The engine ID is also used in combination with user passwords to generate the security keys for authenticating and encrypting SNMPv3 packets.

Command Usage

A local engine ID is automatically generated that is unique to the switch. This is referred to as the default engine ID. If the local engine ID is deleted or changed, all SNMP users will be cleared. You will need to reconfigure all existing users.

Parameters

These parameters are displayed:

- ◆ **Engine ID** – A new engine ID can be specified by entering 9 to 64 hexadecimal characters (5 to 32 octets in hexadecimal format). If an odd number of characters are specified, a trailing zero is added to the value to fill in the last octet. For example, the value “123456789” is equivalent to “1234567890”.
- ◆ **Engine Boots** – The number of times that the engine has (re-)initialized since the SNMP Engine ID was last configured.

Web Interface

To configure the local SNMP engine ID:

1. Click Administration, SNMP.
2. Select Configure Engine from the Step list.
3. Select Set Engine ID from the Action list.
4. Enter an ID of a least 9 hexadecimal characters.
5. Click Apply

Figure 278: Configuring the Local Engine ID for SNMP

Administration > SNMP

Step: 2. Configure Engine Action: Set Engine ID

Engine ID 800001030300000c00000000

Engine Boots 5

Default 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 458.](#))

Parameters

These parameters are displayed:

- ◆ **Remote Engine ID** – The engine ID can be specified by entering 9 to 64 hexadecimal characters (5 to 32 octets in hexadecimal format). If an odd number of characters are specified, a trailing zero is added to the value to fill in the last octet. For example, the value "123456789" is equivalent to "1234567890".
- ◆ **Remote IP Host** – The IPv4 address of a remote management station which is using the specified engine ID.

Web Interface

To configure a remote SNMP engine ID:

1. Click Administration, SNMP.
2. Select Configure Engine from the Step list.
3. Select Add Remote Engine from the Action list.
4. Enter an ID of a least 9 hexadecimal characters, and the IP address of the remote host.
5. Click Apply


Figure 279: Configuring a Remote Engine ID for SNMP

The screenshot shows a web interface titled "Administration > SNMP". Below the title, there are two dropdown menus: "Step:" with "2. Configure Engine" selected, and "Action:" with "Add Remote Engine" selected. Below these, there are two text input fields. The first field is labeled "Remote Engine ID:" and contains the value "5432100000". The second field is labeled "Remote IP Host:" and contains the value "192.168.1.19". At the bottom right of the form, 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 280: Showing Remote Engine IDs for SNMP



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 281: Creating an SNMP View

Administration > SNMP

Step: 3. Configure View Action: Add View

View Name: Entry.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 282: Showing SNMP Views

Administration > SNMP

Step: 3. Configure View Action: Show View

SNMPv3 View List Test 2

	View Name
☐	Entry.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 283: Adding an OID Subtree to an SNMP View

Administration > SNMP

Step: 3. Configure View Action: Add OID Subtree

View Name: rEntry 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 284: Showing the OID Subtree Configured for SNMP Views

Administration > SNMP

Step: 3. Configure View Action: Show OID Subtree

View Name: rEntry a

SNMPv3 View OID Subtree List: Test 3

	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 33: Supported Notification Messages

Model	Level	Group
<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.

Table 33: Supported Notification Messages (Continued)

Model	Level	Group
linkUp*	1.3.6.1.6.3.1.1.5.4	A linkUp trap signifies that the SNMP entity, acting in an agent role, has detected that the ifOperStatus object for one of its communication links left the down state and transitioned into some other state (but not into the notPresent state). This other state is indicated by the included value of ifOperStatus.
authenticationFailure*	1.3.6.1.6.3.1.1.5.5	An authenticationFailure trap signifies that the SNMPv2 entity, acting in an agent role, has received a protocol message that is not properly authenticated. While all implementations of the SNMPv2 must be capable of generating this trap, the snmpEnableAuthenTraps object indicates whether this trap will be generated.
<i>RMON Events (V2)</i>		
risingAlarm	1.3.6.1.2.1.16.0.1	The SNMP trap that is generated when an alarm entry crosses its rising threshold and generates an event that is configured for sending SNMP traps.
fallingAlarm	1.3.6.1.2.1.16.0.2	The SNMP trap that is generated when an alarm entry crosses its falling threshold and generates an event that is configured for sending SNMP traps.
<i>Private Traps</i>		
swPowerStatusChangeTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.1	This trap is sent when the power state changes.
swPortSecurityTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.36	This trap is sent when the port is being intruded. This trap will only be sent when the portSecActionTrap is enabled.
swIpFilterRejectTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.40	This trap is sent when an incorrect IP address is rejected by the IP Filter.
swAtcBcastStormAlarmFireTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.70	When broadcast traffic is detected as a storm, this trap is fired.
swAtcBcastStormAlarmClearTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.71	When a broadcast storm is detected as normal traffic, this trap is fired.
swAtcBcastStormTcApplyTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.72	When ATC is activated, this trap is fired.
swAtcBcastStormTcReleaseTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.73	When ATC is released, this trap is fired.
swAtcMcastStormAlarmFireTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.74	When multicast traffic is detected as the storm, this trap is fired.
swAtcMcastStormAlarmClearTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.75	When multicast storm is detected as normal traffic, this trap is fired.
swAtcMcastStormTcApplyTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.76	When ATC is activated, this trap is fired.
swAtcMcastStormTcReleaseTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.77	When ATC is released, this trap is fired.
stpBpduGuardPortShutdownTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.91	This trap will be sent when an interface is shut down because of BPDU guard.
swLoopbackDetectionTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.95	This trap is sent when loopback BPDUs have been detected.
autoUpgradeTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.104	This trap is sent when auto upgrade is executed.
swCpuUtiRisingNotification	1.3.6.1.4.1.259.10.1.55.102.2.1.0.107	This notification indicates that the CPU utilization has risen from cpuUtiFallingThreshold to cpuUtiRisingThreshold.

Table 33: Supported Notification Messages (Continued)

Model	Level	Group
swCpuUtiFallingNotification	1.3.6.1.4.1.259.10.1.55.102.2.1.0.108	This notification indicates that the CPU utilization has fallen from <code>cpuUtiRisingThreshold</code> to <code>cpuUtiFallingThreshold</code> .
swMemoryUtiRisingThreshold Notification	1.3.6.1.4.1.259.10.1.55.102.2.1.0.109	This notification indicates that the memory utilization has risen from <code>memoryUtiFallingThreshold</code> to <code>memoryUtiRisingThreshold</code> .
swMemoryUtiFallingThreshold Notification	1.3.6.1.4.1.259.10.1.55.102.2.1.0.110	This notification indicates that the memory utilization has fallen from <code>memoryUtiRisingThreshold</code> to <code>memoryUtiFallingThreshold</code> .
dhcpRogueServerAttackTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.114	This trap is sent when receiving a DHCP packet from a rogue server.
macNotificationTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.138	This trap is sent when there are changes of the dynamic MAC addresses on the switch.
lbdDetectionTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.141	This trap is sent when a loopback condition is detected by LBD.
lbdRecoveryTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.142	This trap is sent when a recovery is done by LBD.
sfpThresholdAlarmWarnTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.189	This trap is sent when the sfp's A/D quantity is not within alarm/warning thresholds.
userAuthenticationFailureTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.199	This trap will be triggered if authentication fails.
userAuthenticationSuccessTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.200	This trap will be triggered if authentication is successful.
loginTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.201	This trap is sent when user logs in.
logoutTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.202	This trap is sent when user logs out.
fileCopyTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.208	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	1.3.6.1.4.1.259.10.1.55.102.2.1.0.209	This trap is sent when a user account is created.
userauthDeleteUserTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.210	This trap is sent when a user account is deleted.
userauthModifyUserPrivilegeTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.211	This trap is sent when user privilege is modified.
cpuGuardControlTrap	1.3.6.1.4.1.259.10.1.55.102.2.1.0.213	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	1.3.6.1.4.1.259.10.1.55.102.2.1.0.214	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 285: Creating an SNMP Group

Administration > SNMP

Step: 4. Configure Group Actions: Add

Group Name:

Security Model:

Security Level:

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 286: Showing SNMP Groups

Administration > SNMP

Step: 4. Configure Group Actions: 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

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 287: Configuring Local SNMPv3 Users

Administration > SNMP

Step: 5. Configure User Actions: Add SNMPv3 Local User

SNMPv3 User

User Name:

Group Name: ☐ public ☒ rld

Security Model:

Security Level:

User Authentication

Authentication Protocol:

Authentication Password:

Data Privacy

Privacy Protocol:

Privacy Password:

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 288: Showing Local SNMPv3 Users

Administration > SNMP

Step: 5. Configure User Actions: Show SNMPv3 Local User

SNMPv3 Local User List Total: 1

	User Name	Group Name	Model	Level	Authentication	Privacy
☐	chris	rld	v3	authPriv	MD5	DES56

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 289: Changing a Local SNMPv3 User Group

The screenshot shows a web interface for configuring SNMPv3 users. The breadcrumb is 'Administration > SNMP'. The 'Steps' dropdown is set to '5. Configure User' and the 'Action' dropdown is set to 'Change SNMPv3 Local User Group'. Under 'User Name', there is a dropdown menu with 'chris' selected. Under 'Group Name', there is a text input field containing 'bart' and a dropdown menu with 'public' selected. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 461](#) and [“Specifying a Remote Engine ID” on page 448.](#))

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** – IPv4 address of the remote device where the user resides.
- ◆ **Security Model** – The user security model; SNMP v1, v2c or v3. (Default: v3)
- ◆ **Security Level** – The following security levels are only used for the groups assigned to the SNMP security model:

- **noAuthNoPriv** – There is no authentication or encryption used in SNMP communications. (This is the default security level.)
- **AuthNoPriv** – SNMP communications use authentication, but the data is not encrypted.
- **AuthPriv** – SNMP communications use both authentication and encryption.
- ◆ **Authentication Protocol** – The method used for user authentication. (Options: MD5, SHA; Default: MD5)
- ◆ **Authentication Password** – A minimum of eight plain text characters is required.
- ◆ **Privacy Protocol** – The encryption algorithm 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 290: Configuring Remote SNMPv3 Users

Administration > SNMP

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

SNMPv3 User

User Name:

Group Name: ☐ public ☒ rld

Remote IP:

Security Model:

Security Level:

User Authentication

Authentication Protocol:

Authentication Password:

Data Privacy

Privacy Protocol:

Privacy Password:

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 291: 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
☐	nark	rld	8432100000	v3	authPriv	MD5	DES56

Specifying Trap Managers

Use the Administration > SNMP (Configure Trap) page to specify the host devices to be sent traps and the types of traps to send. Traps indicating status changes are issued by the switch to the specified trap managers. You must specify trap managers so that key events are reported by this switch to your management station (using network management software). You can specify up to five management stations that will receive authentication failure messages and other trap messages from the switch.

Command Usage

- ◆ Notifications are issued by the switch as trap messages by default. The recipient of a trap message does not send a response to the switch. Traps are therefore not as reliable as inform messages, which include a request for acknowledgment 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 445](#)).
2. Create a view with the required notification messages ([page 449](#)).
3. Configure the group (matching the community string specified on the Configure Trap - Add page) to include the required notify view ([page 452](#)).
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 445](#)).
2. Create a remote SNMPv3 user to use in the message exchange process ([page 456](#)). 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 449](#)).
4. Create a group that includes the required notify view ([page 452](#)).
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 456](#)), 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 458](#)), 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 292: Configuring Trap Managers (SNMPv1)



The screenshot shows the 'Administration > SNMP' configuration page. The 'Step' dropdown is set to '6. Configure Trap' and the 'Action' dropdown is set to 'Add'. The form contains the following fields:

Field	Value
IP Address	192.168.0.3
Version	v1
Community String	private
UDP Port (1-65535)	162

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

Figure 293: Configuring Trap Managers (SNMPv2c)



The screenshot shows the 'Administration > SNMP' configuration page for SNMPv2c. The 'Step' dropdown is set to '6. Configure Trap' and the 'Action' dropdown is set to 'Add'. The form contains the following fields:

Field	Value
IP Address	192.168.2.9
Version	v2c
Notification Type	Inform
Timeout (0-2147483647)	
Retry Times (0-255)	
Community String	value
UDP Port (1-65535)	

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

Figure 294: Configuring Trap Managers (SNMPv3)



The screenshot shows the 'Administration > SNMP' configuration page for SNMPv3. The 'Step' dropdown is set to '6. Configure Trap' and the 'Action' dropdown is set to 'Add'. The form contains the following fields:

Field	Value
IP Address	192.168.2.9
Version	v3
Notification Type	Inform
Timeout (0-2147483647)	
Retry Times (0-255)	
Remote User Name	
UDP Port (1-65535)	
Security Level	authPriv

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

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 295: Showing Trap Managers



IP Address	Version	Community String/Trap Name	UDP Port	Security Level	Timeout	Entry Times
192.168.8.4	v3	snmp	162	noAuthNoPriv		
192.168.8.5	v3	bobby	162	noAuthNoPriv		
192.168.8.6	v3	beth	162	authNoPriv		
192.168.8.8	v2c	venice	162		1000	5
192.168.8.8	v3	margaret	162	authPriv	1000	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 81](#). 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 461](#), 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 296: Creating SNMP Notification Logs

The screenshot shows a web interface titled "Administration > SNMP". Below the title, there is a "Step:" dropdown menu set to "7. Configure Notify Filter" and an "Action:" dropdown menu set to "Add". Below these, there are two input fields: "IP Address" with the value "192.168.0.99" and "Filter Profile Name" with the value "R50". 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 297: 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 298: Showing SNMP Statistics



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 299: Configuring an RMON Alarm

The screenshot shows the 'Administration > RMON' configuration page. The 'Step' dropdown is set to '1. Configure Global' and the 'Action' dropdown is set to 'Add'. The 'Alarm' radio button is selected. The configuration fields are as follows:

- Index (1-85535): 1
- Variable: E.1
- Interval (1-31622400): 15 sec
- Sample Type: Delta
- Rising Threshold (0-2147483647): 100
- Rising Event Index (0-85535): 30
- Falling Threshold (0-2147483647): 1
- Falling Event Index (0-85535): 2
- Owner: bill

'Apply' and 'Revert' buttons are at the bottom right.

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 300: Showing Configured RMON Alarms

The screenshot shows the 'Administration > RMON' configuration page with the 'Step' dropdown set to '1. Configure Global' and the 'Action' dropdown set to 'Show'. The 'Alarm' radio button is selected. Below the configuration fields, there is a table titled 'RMON Alarm List' with a total of 38 alarms. The table has the following columns: Index, Status, Variable, Interval, Type, Last Value, Rising Threshold, Rising Event Index, Falling Threshold, Falling Event Index, and Owner.

Index	Status	Variable	Interval	Type	Last Value	Rising Threshold	Rising Event Index	Falling Threshold	Falling Event Index	Owner
1	Valid	1.3.6.1.2.1.10.1.1.1.1	30	Delta	0	802000	0	440400	0	
2	Valid	1.3.6.1.2.1.10.1.1.1.2	30	Delta	0	802000	0	440400	0	
3	Valid	1.3.6.1.2.1.10.1.1.1.3	30	Delta	0	802000	0	440400	0	
4	Valid	1.3.6.1.2.1.10.1.1.1.4	30	Delta	0	802000	0	440400	0	
5	Valid	1.3.6.1.2.1.10.1.1.1.5	30	Delta	0	802000	0	440400	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 411](#)).
 - **Trap** – Sends a trap message to all configured trap managers (see [“Specifying Trap Managers” on page 461](#)).
 - **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 445](#)) prior to configuring it here. (Range: 1-32 characters)
- ◆ **Description** – A comment that describes this event. (Range: 1-127 characters)
- ◆ **Owner** – Name of the person who created this entry. (Range: 1-32 characters)

Web Interface

To configure an RMON event:

1. Click Administration, RMON.
2. Select Configure Global from the Step list.
3. Select Add from the Action list.
4. Click Event.
5. Enter an index number, the type of event to initiate, the community string to send with trap messages, the name of the person who created this event, and a brief description of the event.
6. Click Apply

Figure 301: Configuring an RMON Event

The screenshot shows a web interface titled "Administration > RMON". At the top, there are two dropdown menus: "Step: 1. Configure Global" and "Action: Add". Below these, there are two radio buttons: "Alarm" (unselected) and "Event" (selected). The "Event" configuration form includes the following fields:

- Index (1-85535): 2
- Type: Log and Trap
- Community: private
- Description: for software group
- Owner: david

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

To show configured RMON events:

1. Click Administration, RMON.
2. Select Configure Global from the Step list.
3. Select Show from the Action list.
4. Click Event.

Figure 302: Showing Configured RMON Events

The screenshot shows the 'Administration > RMON' configuration page. At the top, there's a 'Steps' dropdown set to '1. Configure Global' and an 'Action' dropdown set to 'Show'. Below this, there are radio buttons for 'None' and 'Event', with 'Event' selected. The main area is titled 'RMON Event List' and contains a table with 8 columns: Index, Status, Type, Containing, Description, Owner, and Last Filled. The table has 4 rows of data. Below the table are 'Details' and 'Revert' buttons.

Index	Status	Type	Containing	Description	Owner	Last Filled
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 111](#).
- ◆ 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: 1-65535; Default: 50)
The number of buckets granted are displayed on the Show page.
- ◆ **Owner** - Name of the person who created this entry. (Range: 1-32 characters)

Web Interface

To periodically sample statistics on a port:

1. Click Administration, RMON.
2. Select Configure Interface from the Step list.
3. Select Add from the Action list.
4. Click History.
5. Select a port from the list as the data source.
6. Enter an index number, the sampling interval, the number of buckets to use, and the name of the owner for this entry.
7. Click Apply

Figure 303: Configuring an RMON History Sample

The screenshot shows the 'Administration > RMON' web interface. At the top, there is a 'Step' dropdown menu set to '2. Configure Interface' and an 'Action' dropdown menu set to 'Add'. Below these, there are two radio buttons: 'History' (selected) and 'Statistics'. Under the 'History' section, there is a 'Port' dropdown menu set to '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 304: Showing Configured RMON History Samples

Index	Status	Interval	Requested Buckets	Granted Buckets	Owner
1	Valid	1000	0	0	
2	Valid	20	0	0	

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 305: Showing Collected RMON History Samples

History Index	Sample Index	Interval Start	Octets	Packets	Broadcast Packets	Multicast Packets	Undersize Packets	Oversize Packets	Fragments	Jabbers	CRC Align Errors	Collisions	Drop Events	Network Utilization
1	1	00:00:01	758105	3210	81	394	0	0	0	0	0	0	0	0
2	71	00:05:01	21490	78	0	15	0	0	0	0	0	0	0	0
3	72	00:05:31	48821	120	0	15	0	0	0	0	0	0	0	0
4	73	00:06:01	21682	79	0	15	0	0	0	0	0	0	0	0
5	74	00:06:31	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: 1-32 characters)

Web Interface

To enable regular sampling of statistics on a port:

1. Click Administration, RMON.
2. Select Configure Interface from the Step list.
3. Select Add from the Action list.
4. Click Statistics.
5. Select a port from the list as the data source.
6. Enter an index number, and the name of the owner for this entry
7. Click Apply

Figure 306: Configuring an RMON Statistical Sample

Administration > RMON

Step: 2. Configure Interface Action: Add

☐ History ☒ Statistics

Port: 2

Index (1-65535): 130

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 307: 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 308: 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 165](#)).
 2. Add the participating ports to this VLAN (see [“Adding Static Members to VLANs” on page 167](#)), 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 309: Configuring a Switch Cluster

The screenshot shows the 'Administration > Cluster' configuration page. At the top, there's a breadcrumb 'Administration > Cluster' and a 'Step:' dropdown menu set to '1. Configure Global'. Below this, several configuration parameters are listed in a table-like format:

Cluster Status	☒ Enabled
Commander Status	☒ Enabled
IP Pool	10.254.254.1
Role	Commander
Number of Members	2
Number of Candidates	3

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

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 310: Configuring a Cluster Members

The screenshot shows the 'Administration > Cluster' web interface. At the top, there are two dropdown menus: 'Step: 2. Configure Member' and 'Action: Add'. Below these, there are two input fields: 'Member ID (1-36)' and 'MAC Address'. The 'MAC Address' field has a radio button labeled 'Candidate' next to it, and a dropdown menu showing '11-22-33-44-55-11'. There are also 'Apply' and 'Revert' buttons at the bottom right.

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 311: Showing Cluster Members

The screenshot shows the 'Administration > Cluster' web interface with the 'Action' dropdown set to 'Show'. Below the dropdowns, there is a table titled 'Cluster Member List' with 2 rows. The table has columns for 'Member ID', 'Role', 'IP Address', 'MAC Address', and 'Description'. Below the table, there are 'Delete' and 'Revert' buttons.

Member ID	Role	IP Address	MAC Address	Description
1	Active Member	16.254.254.2	70-72-CF-81-1C-82	EC54625-52P
2	Active Member	16.254.254.3	78-12-1F-58-1A-04	EC53518-28T

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 312: Showing Cluster Candidates



The screenshot shows the 'Administration > Cluster' page. At the top, there are two dropdown menus: 'Step: 2. Configure Member' and 'Action: Show Candidate'. Below these is a table titled 'Cluster Candidate List Total: 2'. The table has three columns: 'Role', 'MAC Address', and 'Description'. It contains two rows: one for an 'Active Member' with MAC address '70-72-CF-81-1C-B2' and description 'SC04625-S20', and one for a 'Candidate' with MAC address '70-72-CF-58-2A-D4' and description 'SC53615-267'. A 'Clear' button is located at the bottom right of the table.

Role	MAC Address	Description
Active Member	70-72-CF-81-1C-B2	SC04625-S20
Candidate	70-72-CF-58-2A-D4	SC53615-267

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 313: Managing a Cluster Member



Administration > Cluster					
Step: 3. Show Member					
Cluster Member List Total: 2					
	Member ID	Role	IP Address	MAC Address	Description
☒	1	Active Member	10.254.254.2	78:72:CF:81:1C:82	PCS4620-528
☐	2	Active Member	10.254.254.3	78:72:CF:88:34:04	PCS3610-287
Continue					

Setting a Time Range

Use the Administration > Time Range page to set a time range during which various functions are applied, including applied ACLs or PoE.

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.
- ◆ A maximum of eight rules can be configured for a time range.

Parameters

These parameters are displayed:

Add

- ◆ **Time-Range Name** – Name of a time range. (Range: 1-32 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 Administration, Time Range.
2. Select Add from the Action list.
3. Enter the name of a time range.
4. Click Apply.

Figure 314: Setting the Name of a Time Range



The screenshot shows the 'Administration > Time Range' web interface. At the top, there is a breadcrumb trail 'Administration > Time Range'. Below it, there is a section with 'Action:' and a dropdown menu set to 'Add'. Underneath, there is a 'Time-Range Name' label followed by a text input field containing the text 'rld'. At the bottom right of this section, there are two buttons: 'Apply' and 'Revert'.

To show a list of time ranges:

1. Click Administration, Time Range.
2. Select Show from the Action list.

Figure 315: Showing a List of Time Ranges



The screenshot shows the 'Administration > Time Range' web interface. At the top, there is a breadcrumb trail 'Administration > Time Range'. Below it, there is a section with 'Action:' and a dropdown menu set to 'Show'. Underneath, there is a 'Time-Range List' label followed by a table. The table has a header row with 'Time-Range Name' and a data row with 'rld'. At the bottom right of this section, there are two buttons: 'Create' and 'Revert'.

To configure a rule for a time range:

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

Figure 316: Add a Rule to a Time Range

The screenshot shows the 'Administration > Time Range' configuration page. The 'Action:' dropdown is set to 'Add Rule'. The 'Time-Range' dropdown is set to 'r8d'. The 'Mode' dropdown is set to 'Periodic'. The 'Start' section includes 'Days of the week' set to 'Weekend', 'Hours (0-23)' set to '5', and 'Minutes (0-59)' set to '0'. The 'To' section includes 'Days of the week' set to 'Sunday', 'Hours (0-23)' set to '6', and 'Minutes (0-59)' set to '0'. At the bottom right are 'Apply' and 'Revert' buttons.

To show the rules configured for a time range:

1. Click Administration, Time Range.
2. Select Show Rule from the Action list.

Figure 317: Showing the Rules Configured for a Time Range

The screenshot shows the 'Administration > Time Range' configuration page with the 'Action:' dropdown set to 'Show Rule'. The 'Time-Range' dropdown is set to 'r8d'. Below the form is a table titled 'Time Range Rule List' with a 'Total: 1' indicator. The table has columns for 'Mode', 'Start', and 'End'. The single rule listed is 'Periodic' with 'Start' as 'Weekend 05:00' and 'End' as 'Weekend 06:00'. At the bottom are 'Delete' and 'Revert' buttons.

Mode	Start	End
Periodic	Weekend 05:00	Weekend 06:00

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 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.

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 318: 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 319 on page 489](#) (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 319 on page 489 (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 319: Ring Interconnection Architecture (Multi-ring/Ladder Network)



Configuration Guidelines for ERPS

1. Create an ERPS ring ([Configure Ring – Add](#)): The ring name is used as an index in the G.8032 database. Specify the east and west interfaces. 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.
2. Configure VLAN groups to assign to specific ERPS instances ([Configure VLAN Group – Add](#)).
3. Create ERPS instances ([Configure Instance – Add](#)): Specify an instance name and optionally an ID number.

4. Configure the RPL owner ([Configure Instance – 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.
5. Configure ERPS timers ([Configure Instance – 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.
6. Configure the ERPS Control VLAN ([Configure Instance – 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.
7. Select the Inclusion VLAN groups ([Configure Instance – Configure Details](#)): Select the ERPS VLAN groups to assign to the instances.
8. Enable ERPS ([Configure Global](#)): Before enabling a ring and instance as described in the next steps, first globally enable ERPS on the switch. If ERPS has not yet been enabled or has been disabled, no ERPS rings will work.
9. Enable an ERPS ring ([Configure Ring – Add/Modify](#)): Before an ERPS instance can work, the physical ring must be enabled.
10. Enable an ERPS instance ([Configure Instance – Configure Details](#)): Before an ERPS instance can work, it must be enabled. When configuration is completed and the instance 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 Status on any node.
11. Display ERPS status information ([Configure Instance – Show](#)): Display ERPS status information for all configured rings.

Configuration Limitations for ERPS

The following configuration limitations apply to ERPS:

- ◆ Ring ports cannot 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.

- ◆ One VLAN must be added to an ERPS instance 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 instance (by setting the Status on the [Configure Instance – Configure Details](#) page).

- ◆ **ERPS 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.

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. Set the ERPS Node ID or leave it as the CPU MAC address.
5. Click Apply.

Figure 320: Setting ERPS Global Status

The screenshot shows a web interface for configuring ERPS. The breadcrumb is "Administration > ERPS". Below it is a "Steps" dropdown menu with "1. Configure Global" selected. The main configuration area has two fields: "ERPS Status" with a checked checkbox labeled "Enabled", and "ERPS Node ID" with a text input field containing "8C-EA-1B-0F-CE-F7" and a placeholder "(xx-xx-xx-xx-xx-xx of 100000000000)". At the bottom right are "Apply" and "Revert" buttons.

ERPS VLAN Group Configuration Use the Administration > ERPS (Configure VLAN Group) pages to configure ERPS VLAN groups.

Command Usage

- ◆ A set of VLANs in an Ethernet ring can be grouped into several subsets and applied to an ERPS instance.
- ◆ A VLAN group configuration is allowed to be deleted only if all associations are removed.

Parameters

These parameters are displayed:

- ◆ **VLAN Group Name** – Name of the VLAN group. (Range: 1-12 characters).
- ◆ **VLAN List** – A single VLAN ID, a list of VLAN IDs separated by commas, or a range of VLANs defined by two VLAN IDs separated by a hyphen.

Web Interface

To configure ERPS VLAN groups on the switch:

1. Click Administration, ERPS.
2. Select Configure VLAN Group from the Step list.
3. Select Add from the Action list.
4. Define a VLAN Group name.
5. Configure VLANs to add to the group.
6. Click Apply.

Figure 321: Configuring ERPS VLAN Groups

The screenshot shows a web interface for configuring ERPS. The breadcrumb is 'Administration > ERPS'. Below it, a step indicator shows 'Step: 2. Configure VLAN Group' and an 'Action: Add' dropdown. There are two input fields: 'VLAN Group Name' and 'VLAN List'. At the bottom right are 'Apply' and 'Revert' buttons.

ERPS Ring Configuration Use the Administration > ERPS (Configure Ring) pages to configure a physical ERPS ring.

Command Usage

- ◆ The switch can support ERPS rings up to half the number of physical ports on the switch.
- ◆ 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.
- ◆ Note that a ring port cannot be configured as a member of a spanning tree, a dynamic trunk, or a static trunk.
- ◆ If a port channel (static trunk) is specified as a ring port, it can not be destroyed before it is removed from the ring configuration. A static trunk will be treated as a signal fault, if it contains no member ports or all of its member ports are in signal fault. If a static trunk is configured as a ring port prior to assigning any member ports, spanning tree will be disabled for the first member port assigned to the static trunk.
- ◆ VLANs that are on the exclusion list are **not** protected by the ERPS ring. Any VLAN not listed on either the inclusion or exclusion list will be blocked on ring ports.
- ◆ Traffic from control VLANs, inclusion VLANs, and exclusion VLANs of an ERPS ring will be forwarded by non-ERPS ring ports.

Parameters

These parameters are displayed:

- ◆ **Ring Name** – Name of the VLAN group. (Range: 1-12 characters).
- ◆ **West** – The port that connects to the next ring node to the west.
- ◆ **East** – The port that connects to the next ring node to the east.

- ◆ **Exclusion VLANs** – VLAN groups that are to be on the exclusion list of a physical ERPS ring.
- ◆ **Status** – Enables ERPS on the configured ring. (Default: Disabled)
ERPS must be enabled globally on the switch before it can be enabled on an ERPS ring.

Web Interface

To configure a physical ERPS ring on the switch:

1. Click Administration, ERPS.
2. Select Configure Ring from the Step list.
3. Select Add from the Action list.
4. Define a ring name.
5. Select the West and East ports for ring connections.
6. Select any VLAN groups to exclude from the ERPS ring.
7. Mark the Status check box to enable the ERPS ring.
8. Click Apply.

Figure 322: Configuring an ERPS Ring

The screenshot shows the 'Administration > ERPS' web interface. At the top, there's a breadcrumb 'Administration > ERPS' and a green status icon. Below it, a 'Step:' dropdown is set to '3. Configure Ring' and an 'Action:' dropdown is set to 'Add'. The main form contains the following fields:

- Ring Name:** A text input field.
- West:** A dropdown menu with a selection icon.
- East:** A dropdown menu with a selection icon.
- Exclusion VLANs:** A section with a 'Total 1' indicator. It contains a table with two columns: a checkbox and 'VLAN Group'. One row is visible with the group name 'test-group'.
- Status:** A checkbox labeled 'Enabled'.

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

ERPS Instance Configuration Use the Administration > ERPS (Configure Instance) pages to configure an ERPS instance.

Command Usage

An ERPS instance 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 491](#)) before a ring instance can start

running. Once enabled, the RPL owner node and non-owner node state machines will start, and the ring instance will enter the active state.

Parameters

These parameters are displayed:

Add

- ◆ **Instance Name** – Name of an ERPS ring. (Range: 1-12 characters)
- ◆ **ID** – ERPS instance identifier used in R-APS messages. (Range: 1-255)

Show

- ◆ **Name** – Name of a configured ERPS instance.
- ◆ **ID** – ERPS ring identifier used in R-APS messages.
- ◆ **Status** – Shows whether ERPS is enabled on the switch.
- ◆ **Physical Ring** – Shows the ERPS ring name.
- ◆ **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.
- ◆ **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.

- **Down** – The interface is not linked up.
- **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 MEP used to monitor the status on this link.
- ◆ **RPL** – Shows if this node is connected to the RPL.
- ◆ **Inclusion VLAN Groups** – ERPS VLAN groups configured for this instance.

Configure Details

- ◆ **Instance Name** – Name of a configured ERPS instance. (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.
- ◆ **Instance 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.
- ◆ **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 165](#)), add the ring ports for the east and west interface as tagged members to this VLAN (see [“Adding Static Members to VLANs” on page 167](#)), 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.
- ◆ **Major Ring** – The ERPS instance used for sending control packets.

This switch can support multiple instances. However, ERPS control packets can only be sent on one instance. This parameter is used to indicate that the current ring instance is a secondary ring, and to specify the major ring instance that 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 ring has two physical

ports, this ring can only be a major ring, not a secondary ring (or sub-ring) which can have only one physical ring port. The major ring therefore cannot be set if the east port is already configured.

◆ **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.

◆ **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 canceled 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.
- ◆ 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 323: 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 324: 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.

- ◆ **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: 1-12 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.

- ◆ **Physical Ring** – Specifies the name of the physical ring on which the configured instance runs.
- ◆ **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.
 - **Down** – The interface is not linked up.
 - **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 MEPs used to monitor the link on a ring node.
- ◆ **RPL** – If node is connected to the RPL, this shows by which interface.
- ◆ **Inclusion VLANs** – Specifies VLAN groups that are to be on the inclusion list of an the ERPS instance. VLANs that are on the inclusion list are protected by the ERPS instance. Any VLAN not listed on either the inclusion or exclusion list will be blocked on ring ports.
- ◆ **Status** – Activates the current ERPS instance. (Default: Disabled)

Before enabling a ring instance, the global ERPS function should be enabled see ([“ERPS Global Configuration” on page 491](#)), the physical ring east and west 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.

Web Interface

To create an ERPS instance:

1. Click Administration, ERPS.
2. Select Configure Instance from the Step list.
3. Select Add from the Action list.
4. Enter a name and optional identifier for the instance.
5. Click Apply.

Figure 325: Creating an ERPS Instance



The screenshot shows a web interface titled "Administration > ERPS". At the top, there are two dropdown menus: "Step:" with "4. Configure Instance" selected, and "Action:" with "Add" selected. Below these, there is a form with two fields: "Instance Name" with an empty text box, and "ID (1-255)" with a checkbox and an empty text box. At the bottom right of the form, there are two buttons: "Apply" and "Revert".

To configure the ERPS parameters for an instance:

1. Click Administration, ERPS.
2. Select Configure Instance from the Step list.
3. Select Configure Details from the Action list.
4. Configure the ERPS instance 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 326: Configuring ERPS Instance Details

Administration > ERPS

Step: 4. Configure Instance Action: Configure Details

Instance Name: inst1

Instance ID: 1

Control VLAN: ☒ 1

Node State: init

Node Type: None

Major Ring: ☐ ERPSinst2

Version: 2

REG Level (0-7): 1

Reversion: ☒ Enabled

R-APS with VC: ☒ Enabled

R-APS Del MAC: ☒ Enabled

Propagate TC: ☐ Enabled

Holdoff Timer (0-10000): 0 ms

Guard Timer (10-2000): 500 ms

WTS Timer: 5500 ms

WTR Timer (1-12): 5 min

WTS Expire:

WTR Expire:

Physical Ring: ☒ Inst-req

West		East	
Interface	Port State	Interface	Port State
	Unknown		Unknown
Local SF	No	Local SF	No
Local FS	No	Local FS	No
Local MS	No	Local MS	No
MEP		MEP	
RPL	No	RPL	

Inclusion VLANs: Total 5

	VLAN Group
☒	Inst-group
☐	ERPSgroup2
☐	ERPSgroup3
☐	ERPSgroup4
☐	ERPSgroup5

Status: ☐ Enabled

Apply Reset

To show the configured ERPS instances:

1. Click Administration, ERPS.
2. Select Configure Instance from the Step list.
3. Select Show from the Action list.

Figure 327: Showing Configured ERPS Instances

The screenshot shows the 'Administration > ERPS' configuration page. At the top, there's a 'Step: 4. Configure Instance' and an 'Action: Show' dropdown. Below this is the 'Instance List' table. The table has columns: Name, ID, Status, Physical Ring, Control VLAN, Node Name, Type, Wt, Interface, Port State, Local ST, Local FS, Local MS, RPL, RPL, and Inclusion VLAN Group. There are three instances listed: 'Inst1' (ID 1, Disabled, Physical Ring: Inst-req, Control VLAN: 1, Node Name: Inst1, Type: Ring, Wt: West, Interface: Unknown, Port State: Unknown, Local ST: No, Local FS: No, Local MS: No, RPL: No, RPL: No, Inclusion VLAN Group: 1), 'ERPSInst2' (ID 12, Disabled, Physical Ring: , Control VLAN: , Node Name: Inst1, Type: Ring, Wt: West, Interface: Unknown, Port State: Unknown, Local ST: No, Local FS: No, Local MS: No, RPL: No, RPL: No, Inclusion VLAN Group:), and 'ERPSInst3' (ID 13, Disabled, Physical Ring: , Control VLAN: , Node Name: Inst1, Type: Ring, Wt: East, Interface: Unknown, Port State: Unknown, Local ST: No, Local FS: No, Local MS: No, RPL: No, RPL: No, Inclusion VLAN Group:). At the bottom of the table are 'Delete' and 'Reset' buttons.

Name	ID	Status	Physical Ring	Control VLAN	Node Name	Type	Wt	Interface	Port State	Local ST	Local FS	Local MS	RPL	RPL	Inclusion VLAN Group
Inst1	1	Disabled	Inst-req	1	Inst1	Ring	West	Unknown	Unknown	No	No	No	No	No	1
ERPSInst2	12	Disabled			Inst1	Ring	West	Unknown	Unknown	No	No	No	No	No	
ERPSInst3	13	Disabled			Inst1	Ring	East	Unknown	Unknown	No	No	No	No	No	

ERPS Forced and Manual Mode Operations

Use the Administration > ERPS (Configure Instance - Configure Operation) page to block a ring port using Forced Switch or Manual Switch commands.

Parameters

These parameters are displayed:

- ◆ **Instance Name** – Name of a configured ERPS instance.
- ◆ **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 34 on page 510](#)). 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 34: 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 34 on page 510](#)). 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 Instance Configuration” on page 494.](#))

Web Interface

To block a ring port:

1. Click Administration, ERPS.
2. Select Configure Instance from the Step list.
3. Select Configure Operation from the Action list.
4. Select the instance name from the drop-down list.

5. Specify a Forced Switch, Manual Switch, or Clear operation.
6. Click Apply.

Figure 328: Blocking an ERPS Ring Port

Administration > ERPS

Step: 4. Configure Instance Action: Configure Operation

Instance Name: Inst1


Operation: Forced-switch West

Apply

MLAG Configuration

A multi-chassis link aggregation group (MLAG) is a pair of links that terminate on two cooperating switches and appear as an ordinary link aggregation group (LAG). The cooperating switches are MLAG peer switches and communicate through an interface called a peer link. While the peer link's primary purpose is exchanging MLAG control information between peer switches, but also carries data traffic from devices that are attached to only one MLAG peer and have no alternative path. An MLAG domain consists of the peer switches and the control links that connect these switches.

Figure 329: MLAG Domain Topology



Usage Guidelines

- ◆ MLAG must be enabled globally from the Configure Global page.
- ◆ The MLAG domain ID and peer link must be set using the Configure Domain - Add page.
- ◆ The MLAG ID, associated MLAG domain ID and MLAG member must be configured using the Configure Group - Add page. The associated MLAG domain may be nonexistent, which causes MLAG to be inactive locally.
- ◆ For a port to be configured as MLAG peer link or member:
 - STP status of the port must be disabled.
 - LACP status of the port must be disabled.
 - The port must not be any type of traffic segmentation port.

MLAG Restrictions

- ◆ Traffic segmentation up-link/down-link port cannot be configured on an MLAG member or peer link.
- ◆ All actions that cause a port to become nonexistent, such as deleting a trunk port, adding a port to a trunk, or enabling LACP, are not allowed for an MLAG member or peer link. Also, a trunk member port is not allowed to be an MLAG member or peer link.
- ◆ STP cannot be enabled on a peer link or an MLAG member. An STP enabled port cannot be configured as a peer link or an MLAG member.

MLAG Global Configuration Use the Administration > MLAG > Global page to enable MLAG functionality on the switch.

Parameters

These parameters are displayed:

- ◆ **MLAG Status** – Enables or disables MLAG functions. (Default: Enabled)

Web Interface

To enable MLAG functionality on the switch:

1. Click Administration, MLAG.
2. Select Configure Global from the Step list
3. Set the MLAG administrative status to enabled.
4. Click Apply.

Figure 330: Configuring MLAG Global Status



MLAG Domain Configuration Use the Administration > MLAG > Domain page to configure MLAG domains and peer links on the switch.

Usage Guidelines

- ◆ There shall be one and only one peer link for a pair of MLAG devices in the same MLAG domain. (See [Figure 329](#).)
- ◆ The peer link can be a normal port or a static trunk.
- ◆ MAC learning is automatically disabled for the peer link.
- ◆ An MLAG domain is active if the domain ID and a peer link are set.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – Domain identifier. (Range: 1-16 characters)
- ◆ **Interface** – The peer link port or trunk identifier.

Web Interface

To enable MLAG functionality on the switch:

1. Click Administration, MLAG.
2. Select Configure Domain from the Step list
3. Set an MLAG domain identifier.
4. Select one port or trunk interface as the peer link to another MLAG device.
5. Click Apply.

Figure 331: Configuring MLAG Domains

Administration > MLAG

Step: 3. Configure Domain Action: Add

Domain ID:

Interface: ☒ Port ☐ Trunk

Apply Revert


MLAG Group Configuration Use the Administration > MLAG > Group page to configure MLAG member interfaces on the switch.

Usage Guidelines

- ◆ An MLAG domain can have two and only two MLAG devices. (See [Figure 329](#).)
- ◆ An MLAG domain may have many MLAGs.
- ◆ An MLAG can belong to one and only one MLAG domain.
- ◆ The associated MLAG domain may be nonexistent, which causes the MLAG to be inactive locally.
- ◆ There can be one and only one MLAG member for each MLAG on an MLAG device.
- ◆ The MLAG member can be a normal port or a static trunk.
- ◆ An MLAG member is active if the MLAG ID is set and the associated MLAG domain is active.
- ◆ An MLAG is formed when the peer MLAG members are both active.
- ◆ The following items apply when an MLAG is formed.
 - When an MLAG member is operationally up and the MLAG peer member is not operationally down, all traffic from the peer link can not be forwarded to the MLAG member.
 - When an MLAG member is operationally up and the MLAG peer member is operationally down, all traffic from the peer link can be forwarded to the MLAG member.
 - When an MLAG member is operationally up, all updates for learned MAC addresses on the MLAG peer member will be synced to the MLAG member automatically.

- When an MLAG member is operationally down, all updates for learned MAC addresses on the MLAG peer member will be synced through the peer link automatically.

Figure 332: MLAG Peer Operation



- ◆ When the MLAG peer member is down or nonexistent, learned MAC addresses are synced through the peer link for the MLAG will be removed automatically.

Parameters

These parameters are displayed:

- ◆ **MLAG ID** – MLAG identifier. (Range: 1-1000)
- ◆ **Domain ID** – Domain identifier. (Range: 1-16 characters)
- ◆ **Interface** – The MLAG member port or trunk identifier.

Web Interface

To enable MLAG functionality on the switch:

1. Click Administration, MLAG.
2. Select Configure Group from the Step list.
3. Set an MLAG identifier.
4. Set the MLAG domain identifier.
5. Select a port or trunk interface as an MLAG member.
6. Click Apply.

Figure 333: Configuring MLAG Groups

The screenshot shows a web-based configuration interface for MLAG. At the top, it says "Administration > MLAG". Below this, there's a "Steps:" dropdown menu set to "3. Configure Group" and an "Action:" dropdown menu set to "Add". The main configuration area has three input fields: "MLAG ID", "Domain ID", and "Interface". The "Interface" field has two radio buttons: "Port" (which is selected) and "Trunk". Below these fields are two buttons: "Apply" and "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 334: 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 335: 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 enabled¹³ 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 enabled¹³, and a CCM is received from all remote MEPs configured in the static list for this maintenance association¹⁴.
- ◆ **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.

13. Cross-checking must be enabled for this type of trap to be reported (see ["Configuring CFM Maintenance Associations"](#)).

14. See ["Configuring Maintenance End Points"](#).

A MEP Missing trap is sent if cross-checking is enabled¹³, and no CCM is received for a remote MEP configured in the static list¹⁴.

◆ **Cross Check MEP Unknown** – Sends a trap if an unconfigured MEP comes up.

A MEP Unknown trap is sent if cross-checking is enabled¹³, and a CCM is received from a remote MEP that is not configured in the static list¹⁴.

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 336: Configuring Global Settings for CFM

The screenshot displays the 'Administration > CFM' web interface. At the top, a 'Step' dropdown menu is set to '1. Configure Global'. Below this, the 'Global Configuration' section contains the following settings:

- 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

The 'SNMP Trap Configuration' section contains the following settings:

- Connectivity Check Config: ☐ Enabled
- Connectivity Check Loop: ☐ Enabled
- Connectivity Check MEP Down: ☐ Enabled
- Connectivity Check MEP Up: ☐ Enabled
- Cross Check MA Up: ☐ Enabled
- Cross Check MEP Missing: ☐ Enabled
- Cross Check MEP Unknown: ☐ Enabled

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

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 337: Configuring Interfaces for CFM



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 35: 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 36: 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 338: Configuring Maintenance Domains

Administration > CFM

Step: 1. Configure MD Action: Add

MD Index (1-65536): 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 339: 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 340: 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 text input field:

- MD Index: 1
- MEP Archive Hold Time (1-65535): 100
- MEP Fault Notify Lowest Priority (1-6): 2
- MEP Fault Notify Alarm Time (3-18): 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 534](#)).
- ◆ An MA must be defined before any associated DSAPs or remote MEPs can be assigned (see [“Configuring Remote Maintenance End Points” on page 536](#)).
- ◆ Multiple domains at the same maintenance level cannot have an MA on the same VLAN (see [“Configuring CFM Maintenance Domains” on page 525](#)).
- ◆ Before removing an MA, first remove the MEPs assigned to it (see [“Configuring Maintenance End Points” on page 534](#)).
- ◆ For a detailed description of the MIP types, refer to the Command Usage section under [“Configuring CFM Maintenance Domains” on page 525](#).

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: 1-43¹⁵ 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)

15. 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: Enabled)
- ◆ **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. (Default: Enabled)

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: Enabled)
- ◆ **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 341: Creating Maintenance Associations

Administration > CFM

Steps: 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 342: Showing Maintenance Associations

Administration > CFM

Steps: 4. Configure MA Action: Show

MD Index: 1

CFM MA List	MD Index	MA Name	Primary VLAN	MIP Creation Type
1	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 343: Configuring Detailed Settings for Maintenance Associations

The screenshot shows a web-based configuration interface for 'Administration > CFM'. At the top, there is a breadcrumb trail '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 several parameters:

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 configuration area, there are two buttons: 'Apply' and 'Revert'.

Configuring Maintenance End Points

Use the Administration > CFM (Configure MEP – Add) page to configure Maintenance End Points (MEPs). MEPs, also called Domain Service Access Points (DSAPs), must be configured at the domain boundary to provide management access for each maintenance association.

Command Usage

- ◆ CFM elements must be configured in the following order: (1) maintenance domain at the same level as the MEP to be configured (see ["Configuring CFM Maintenance Domains"](#)), (2) maintenance association within the domain (see ["Configuring CFM Maintenance Associations"](#)), and (3) finally the MEPs using the MEP List.
- ◆ An interface may belong to more than one domain, or to different MAs in different domains.
- ◆ To change the MEP's MA or the direction it faces, first delete the MEP, and then create a new one.

Parameters

These parameters are displayed:

- ◆ **MD Index** – Domain index. (Range: 1-65535)

- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)
- ◆ **MEP ID** – Maintenance end point identifier. (Range: 1-8191)
- ◆ **MEP Direction** – Up indicates that the MEP faces inward toward the switch cross-connect matrix, and transmits CFM messages towards, and receives them from, the direction of the internal bridge relay mechanism. If the **Up** option is not selected, then the MEP is facing away from the switch, and transmits CFM messages towards, and receives them from, the direction of the physical medium.
- ◆ **Interface** – Indicates a port or trunk.

Web Interface

To configure a maintenance end point:

1. Click Administration, CFM.
2. Select Configure MEP from the Step list.
3. Select Add from the Action list.
4. Select an entry from MD Index and MA Index.
5. Specify the MEPs assigned to each MA, set the MEP identifier, the direction in which the MEP faces, and the physical interface serving as the DSAP.
6. Click Apply.

Figure 344: Configuring Maintenance End Points

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 345: Showing Maintenance End Points



Configuring Remote Maintenance End Points

Use the Administration > CFM (Configure Remote MEP – Add) page to specify remote maintenance end points (MEPs) set on other CFM-enabled devices within a common MA. Remote MEPs can be added to a static list in this manner to verify that each entry has been properly configured and is operational. When cross-checking is enabled, the list of statically configured remote MEPs is compared against the MEPs learned through continuity check messages (CCMs), and any discrepancies reported via SNMP traps.

Command Usage

- ◆ All MEPs that exist on other devices inside a maintenance association should be statically configured to ensure full connectivity through the cross-check process.
- ◆ Remote MEPs can only be configured if local domain service access points (DSAPs) have already been created (see ["Configuring Maintenance End Points"](#)) at the same maintenance level and in the same MA. DSAPs are MEPs that exist on the edge of the domain, and act as primary service access points for end-to-end cross-check, loop-back, and link-trace functions.
- ◆ The MEP cross-check start delay which sets the maximum delay that a device waits for remote MEPs to come up before starting the cross-check operation can be configured on the Configure Global page (see ["Configuring Global Settings for CFM"](#)).
- ◆ SNMP traps for continuity check events discovered by cross-check operations can also be configured on the Configure Global page (see ["Configuring Global Settings for CFM"](#)).

Parameters

These parameters are displayed:

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)

- ◆ **MEP ID** – Identifier for a maintenance end point which exists on another CFM-enabled device within the same MA. (Range: 1-8191)

Web Interface

To configure a remote maintenance end point:

1. Click Administration, CFM.
2. Select Configure Remote MEP from the Step list.
3. Select Add from the Action list.
4. Select an entry from MD Index and MA Index.
5. Specify the remote MEPs which exist on other devices within the same MA.
6. Click Apply.

Figure 346: Configuring Remote Maintenance End Points

The screenshot shows a web interface titled "Administration > CFM". At the top, there are two dropdown menus: "Step:" with the value "4. Configure Remote MEP" and "Action:" with the value "Add". Below these, there are three input fields: "MD Index" with a dropdown menu showing "1", "MA Index" with a dropdown menu showing "1", and "MEP ID (1-8191)" with a text input field containing "2". At the bottom right, there are two buttons: "Apply" and "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 347: Showing Remote Maintenance End Points



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 348: Transmitting Link Trace Messages

Administration > CFM

Step: 7. Transmit Link Trace

MD Index: 1

MA Index: 1

Source MEP ID (1-8191): 1

Target: ☒ MEP ID (1-8191) 2 ☐ MAC Address

TTL (0-255): 5

Apply Revert

Transmitting Loop Back Messages

Use the Administration > CFM (Transmit Loopback) page to transmit Loopback Messages (LBMs). These messages can be used to isolate or verify connectivity faults by submitting a request to a target node (i.e., a remote MEP or MIP) to echo the message back to the source.

Command Usage

- ◆ Loopback messages can be used for fault verification and isolation after automatic detection of a fault or receipt of some other error report. Loopback messages can also be used to confirm the successful restoration or initiation of connectivity. The receiving maintenance point should respond to the loop back message with a loopback reply.
- ◆ The point from which the loopback message is transmitted (i.e., a local DSAP) and the target maintenance point must be within the same MA.
- ◆ If the continuity check database does not have an entry for the specified maintenance point, an error message will be displayed.
- ◆ When using the command line or web interface, the source MEP used by to send a loopback message is chosen by the CFM protocol. However, when using SNMP, the source MEP can be specified by the user.

Parameters

These parameters are displayed:

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)
- ◆ **Source MEP ID** – The identifier of a source MEP that will send the loopback message. (Range: 1-8191)

◆ **Target**

- **MEP ID** – The identifier of a remote MEP that is the target of a loopback message. (Range: 1-8191)
- **MAC Address** – MAC address of a remote MEP that is the target of a loopback message. This address can be entered in either of the following formats: xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx

◆ **Count** – The number of times the loopback message is sent. (Range: 1-1024)

◆ **Packet Size** – The size of the loopback message. (Range: 64-1518 bytes; Default: 64 bytes)

Web Interface

To transmit loopback messages:

1. Click Administration, CFM.
2. Select Transmit Loopback from the Step list.
3. Select an entry from MD Index and MA Index.
4. Specify the source MEP, the target MEP using either its MEP identifier or MAC address, set the number of times the loopback message is to be sent.
5. Click Apply.

Figure 349: Transmitting Loopback Messages

The screenshot shows the 'Administration > CFM' web interface. The 'Step' dropdown is set to '6. Transmit Loopback'. The form contains the following fields:

- MD Index:** A dropdown menu with '1' selected.
- MA Index:** A dropdown menu with '1' selected.
- Source MEP ID (1-8191):** An empty text input field.
- Target:** Two radio buttons are present: 'MEP ID (1-8191)' (selected) and 'MAC Address'. The 'MEP ID' radio button is followed by an empty text input field.
- Counts (1-1024):** An empty text input field.
- Packet Size (64-1518):** An empty text input field followed by the word 'byte'.

A 'Tag' button is located below the form fields. Below the form is a 'Result' section with a text area containing the following output:

```

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 sec. 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 350: Transmitting Delay-Measure Messages

Administration > CFM

Step: 1. 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 6 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 351: Showing Information on Local MEPs



MEP ID	MD Name	Level	Direction	Primary VLAN	Interface	CC Status	MAC Address
1	vlag	3	Down	1	Unit 1 / Port 1	Enabled	06-80-0C-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 352: 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	rd1
MA Name	ma1
MA Name Format	Character String
Level	0
Direction	Up
Interface	Unit 1 / Port 2
CC Status	Enabled
MAC Address	00-1B-05-50-01-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 353: Showing Information on Local MIPs



MD Name	Level	MA Name	Primary VLAN	Interface
vrip	3	rt	1	Unit 1 / Port 2
vrip	3	rt	1	Unit 1 / Port 3
vrip	3	rt	1	Unit 1 / Port 4
vrip	3	rt	1	Unit 1 / Port 5
vrip	3	rt	1	Unit 1 / Port 6
vrip	3	rt	1	Unit 1 / Port 7
vrip	3	rt	1	Unit 1 / Port 8
vrip	3	rt	1	Unit 1 / Port 9
vrip	3	rt	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 354: Showing Information on Remote MEPs

MEP ID	MA Name	Level	Primary VLAN	MEP Up	Remote MAC Address
1	net	0	2	Yes	00:01:02:83:81:65

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 355: Showing Detailed Information on Remote MEPs

The screenshot displays the 'Administration > CFM' web interface. At the top, there are two dropdown menus: 'Step: 10. Show Information' and 'Action: Show Remote MEP Details'. Below these, there are three dropdown menus for selection: 'MD Index' (set to 1), 'MA Index' (set to 1), and 'MEP ID' (set to 1). A 'Query' button is located to the right of these dropdowns. The main area of the interface displays a list of parameters and their corresponding values for the selected MEP:

MD Name	red1
MA Name	ma1
Level	0
MAC Address	00-1B-05-50-91-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	138 / 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 356: Showing the Link Trace Cache



Index	MA	IP/Mac Address	Forwarded	Ingress MAC Address	Egress MAC Address	Ingress Action	Egress Action	Reply
1	device	11-11-11-11-11-11	Forwarded	11-11-11-11-11-12	11-11-11-11-11-13	OK	OK	FDB
2	session	AA-BB-CC-DD-EE-FF	Not Forwarded					

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 alarm¹⁶.
- ◆ **Alarm Time** – The time a defect must exist before a fault alarm is issued¹⁶.
- ◆ **Reset Time** – The time after a fault alarm has been issued, and no defect exists, before another fault alarm can be issued¹⁶.

16. See [“Configuring CFM Maintenance Domains”](#) on page 525.

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 357: Showing Settings for the Fault Notification Generator

MEP ID	MD Name	MA Name	Highest Defect	Lowest Alarm	Alarm Time (sec)	Reset Time (sec)
1	ge	10	NONE	allDef	3	1

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 list¹⁷, 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 list¹⁷, 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.

17. 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 358: Showing Continuity Check Errors



Level	Primary VLAN	MEPID	Interface	Remote MAC	Reason	MA Name
5	2	4E	Unit 1 / Port 18	3C-E1-32-63-04-65	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: 1-28/54)

- ◆ **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 37: 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 359: Enabling OAM for Local Ports

Administration > OAM > Interface

OAM Port List Total: 26

Port	Admin Status	Operational State	Mode	Critical Link Event		Errored Frame		
				Dying Gasp	Critical Event	Status	Window Size (10-65535, 1/10 sec)	thr. Threshold Count (1-65535)
1	☐ Enabled	Disabled	Active ▾	☒ Enabled	☒ Enabled	☒ Enabled	10	1
2	☐ Enabled	Disabled	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

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: 1-28/54)
- ◆ **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 360: Displaying Statistics for OAM Messages

Administration > OAM > Counters

OAM Port Counters Test 28

123

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 361: 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: 1-28/54)
- ◆ **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 362: Displaying Status of Remote Interfaces

Administration > OAM > Remote Interface

OAM Remote Port List: Total 25

Port	MAC Address	OUI	Remote Loopback	Unidirectional Function	Link Monitor	MIB Variable Retrieval
1	30-00-03-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
2	30-00-03-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
3	30-00-03-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
4	30-00-03-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
5	30-00-03-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
6	30-00-03-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
7	30-00-03-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
8	30-00-03-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
9	30-00-03-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
10	30-00-03-00-00-00	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: 1-28/54)
- ◆ **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 38: 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 363: Running a Remote Loop Back Test

The screenshot shows the 'Administration > OAM > Remote Loopback' web interface. At the top, the 'Action' dropdown is set to 'Remote Loopback Test'. Below this, there are configuration fields: 'Port' (a dropdown menu), 'Loopback Mode' (a checkbox labeled 'Enabled' which is checked), and 'Loopback Status' (a text field showing 'Remote Loopback:'). There are 'Apply' and 'Revert' buttons. Below the configuration section, there are input fields for 'Packet Number (1-9999999)' and 'Packet Size (64-1536)' with a 'bytes' label. There are 'Test' and 'Test Result' buttons. At the bottom, a 'Result' section displays the text: '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: 1-28/54)
- ◆ **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 364: Displaying the Results of Remote Loop Back Testing



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 a 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 Mfast (7 seconds).

If the link is instead deemed bidirectional, the curve will use Mfast for the first four subsequent message transmissions and then transition to an Mslow value for all other steady-state transmissions. Mslow is the value configured by this command.
- ◆ **Detection Interval** – Sets the amount of time the switch remains in detection state after discovering a neighbor. (Range: 5-255 seconds; Default: 5 seconds)

When a neighbor device is discovered by UDLD, the switch enters “detection state” and remains in this state for specified detection-interval. After the detection-interval expires, the switch tries to decide whether or the link is unidirectional based on the information collected during the “detection state.”
- ◆ **Recovery Status** – Configures the switch to automatically recover from UDLD disabled port state after a period specified by the Recovery Interval. (Default: Disabled)

When automatic recovery state is changed, any ports shut down by UDLD will be reset.
- ◆ **Recovery Interval** – Specifies the period after which to automatically recover from UDLD disabled port state. (Range: 30-86400 seconds; Default: 300 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 365: Configuring UDLD Protocol Intervals

The screenshot shows the 'Administration > UDLD' configuration page. At the top, there is a breadcrumb trail 'Administration > UDLD'. Below it is a 'Step:' dropdown menu set to '1. Configure Global'. The main configuration area contains four rows of settings: 'Message Interval (7-98)' with a value of 15 seconds, 'Detection Interval (5-255)' with a value of 5 seconds, 'Recovery Status' with a checkbox labeled 'Enabled' that is checked, and 'Recovery Interval (30-35400)' with a value of 100 seconds. At the bottom right of the form are two buttons: 'Apply' and 'Revert'.

Configuring UDLD Interface Settings

Use the Administration > UDLD (Configure Interface) page to enable UDLD and aggressive mode which reduces the shut-down delay after loss of bidirectional connectivity is detected.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28/54)
- ◆ **UDLD** – Enables UDLD on a port. (Default: Disabled)
 - UDLD requires that all the devices connected to the same LAN segment be running the protocol in order for a potential mis-configuration to be detected and for prompt corrective action to be taken.
 - Whenever a UDLD device learns about a new neighbor or receives a resynchronization request from an out-of-synch neighbor, it (re)starts the detection process on its side of the connection and sends N echo messages in reply. (This mechanism implicitly assumes that N packets are sufficient to

get through a link and reach the other end, even though some of them might get dropped during the transmission.)

Since this behavior must be the same on all the neighbors, the sender of the echoes expects to receive an echo in reply. If the detection process ends without the proper echo information being received, the link is considered to be unidirectional.

- ◆ **Aggressive Mode** – Reduces the shut-down delay after loss of bidirectional connectivity is detected. (Default: Disabled)

UDLD can function in two modes: normal mode and aggressive mode.

- In normal mode, determination of link status at the end of the detection process is always based on information received in UDLD messages: whether that's information about the exchange of proper neighbor identification or the absence of such. Hence, albeit bound by a timer, normal mode determinations are always based on gleaned information, and as such are "event-based." If no such information can be obtained (e.g., because of a bidirectional loss of connectivity), UDLD follows a conservative approach to minimize false positives during the detection process and deems a port to be in "undetermined" state. In other words, normal mode will shut down a port only if it can explicitly determine that the associated link is faulty for an extended period of time.
- In aggressive mode, UDLD will also shut down a port if it loses bidirectional connectivity with the neighbor for the same extended period of time (as that mentioned above for normal mode) and subsequently fails repeated last-resort attempts to re-establish communication with the other end of the link. This mode of operation assumes that loss of communication with the neighbor is a meaningful network event in itself, and a symptom of a serious connectivity problem. Because this type of detection can be event-less, and lack of information cannot always be associated to an actual malfunction of the link, this mode is recommended only in certain scenarios (typically only on point-to-point links where no communication failure between two neighbors is admissible).

- ◆ **Operation State** – Shows the UDLD operational state (Disabled, Link down, Link up, Advertisement, Detection, Disabled port, Advertisement - Single neighbor, Advertisement - Multiple neighbors)

- ◆ **Port State** – Shows the UDLD port state (Unknown, Bidirectional, Unidirectional, Transmit-to-receive loop, Mismatch with neighbor state reported, Neighbor's echo is empty)

The state is Unknown if the link is down or not connected to a UDLD-capable device. The state is Bidirectional if the link has a normal two-way connection to a UDLD-capable device. All other states indicate mis-wiring.

- ◆ **Message Interval** – The interval between UDLD probe messages used for the indicated operational state.

- ◆ **Detection Interval** – The period the switch remains in detection state after discovering a neighbor.

Web Interface

To enable UDLD and aggressive mode:

1. Click Administration, UDLD, Configure Interface.
2. Enable UDLD and aggressive mode on the required ports.
3. Click Apply.

Figure 366: 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: 1-28/52)
- ◆ **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, neighborsEchoIsEmpty, 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 367: Displaying UDLD Neighbor Information



Entry	Device ID	Port ID	Device Name	Neighbor State	Expires (seconds)	Message Interval (seconds)	Detection Interval (seconds)
1	8120466	Et-1/1	ICS4100-02P	Bidirectional	24	7	1

LBD Configuration

The switch can be configured to detect general loopback conditions caused by hardware problems or faulty protocol settings. When loopback detection (LBD) is 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.
- ◆ 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 Global Settings for LBD

Use the Administration > LBD (Configure Global) page to enable loopback detection globally, specify the interval at which to transmit control frames, the interval to wait before releasing an interface from shutdown state, the response to a detected loopback, and the traps to send.

Parameters

These parameters are displayed:

- ◆ **Global Status** – Enables loopback detection globally on the switch. (Default: Enabled)
- ◆ **Transmit Interval** – Specifies the interval at which to transmit loopback detection control frames. (Range: 1-32767 seconds; Default: 10 seconds)
- ◆ **Recover Time** – Specifies the interval to wait before the switch automatically releases an interface from shutdown state. (Range: 60-1,000,000 seconds; Default: 60 seconds)

If the recover time is not enabled (checkbox unmarked), all ports placed in shutdown state can be restored to operation using the Release button. To restore a specific port, re-enable Admin status on the Configure Interface page.

The recover-time is the maximum time when recovery is triggered after a loop is detected. The actual interval between recovery and detection will be less than or equal to the recover-time.
- ◆ **Action** – Specifies the protective action the switch takes when a loopback condition is detected. (Options: Block, None, Shutdown; Default: Shutdown)
 - **Block** – When the response to a detected loopback condition is set to block user traffic, and a loopback is detected on a port which a member of a specific VLAN, packets belonging to that VLAN are dropped at the offending port. Under these conditions, loopback detection control frames may be untagged or tagged depending on the port's VLAN membership type. Ingress filtering for the port is enabled automatically if not already enabled by other commands. The port's original setting for ingress filtering will be restored when loopback detection is disabled.
 - **None** - No action is taken.
 - **Shutdown** – When the response to a detected loopback condition is set to shut down a port, and a port receives a control frame sent by itself, this means that the port is in looped state, and the VLAN in the frame payload is also in looped state with the wrong VLAN tag. The looped port is therefore shut down.

When the loopback detection response is changed, any ports placed in shutdown state by the loopback detection process will be immediately restored to operation regardless of the remaining recover time.

- ◆ **Trap** – Sends a trap when a loopback condition is detected, or when the switch recovers from a loopback condition. (Options: Both, Detect, None, Recover; Default: None)
 - **Both** – Sends an SNMP trap message when a loopback condition is detected, or when the switch recovers from a loopback condition.
 - **Detect** – Sends an SNMP trap message when a loopback condition is detected.
 - **None** – Does not send an SNMP trap for loopback detection or recovery.
 - **Recover** – Sends an SNMP trap message when the switch recovers from a loopback condition.
- ◆ **Release** – Releases all interfaces currently shut down by the loopback detection feature.

Web Interface

To configure global settings for LBD:

1. Click Administration, LBD, Configure Global.
2. Make the required configuration changes.
3. Click Apply.

Figure 368: Configuring Global Settings for LBD

The screenshot shows the 'Administration > LBD' configuration page. The 'Step' dropdown is set to '1. Configure Global'. The configuration fields are as follows:

Global Status	☐ Enabled
Transmit Interval (1-32767)	10 sec
☒ Recover Time (60-1066000)	60 sec
Action	Shutdown
Trap	None

At the bottom right, there are 'Apply' and 'Reset' buttons. At the bottom left, there is a 'Release' button with a tooltip that says 'Click this button to release all looped ports manually'.

Configuring Interface Settings for LBD Use the Administration > LBD (Configure Interface) page to enable loopback detection on an interface, to display the loopback operational state, and the VLANs that are looped back.

Parameters

These parameters are displayed:


- ◆ **Interface** – Displays a list of ports or trunks.
 - **Port** – Port Identifier. (Range: 1-28/54)
 - **Trunk** – Trunk Identifier. (Range: 1-28)
- ◆ **Admin State** – Manually enables or disables an interface. (Default: Enabled)
- ◆ **Operation State** – Valid states include Normal or Looped.
- ◆ **Looped VLAN** – Shows the VLANs that are in a looped state.

Web Interface

To configure interface settings for LBD:

1. Click Administration, LBD, Configure Interface.
2. Make the required configuration changes.
3. Click Apply.

Figure 369: Configuring Interface Settings for LBD



Port	Admin State	Operation State	Looped VLAN
1	✓ Enabled	Normal	None
2	✓ Enabled	Normal	None
3	✓ Enabled	Normal	None
4	✓ Enabled	Normal	None
5	✓ Enabled	Normal	None

Smart Pair Configuration

A Smart Pair consists of two ports which are paired to provide layer 2 link redundancy. The pair consists of a primary port and a backup port. All traffic is forwarded through the primary port and the backup port will be set to standby. If the primary port link goes down, the backup port is activated and all traffic is

forwarded through it. If the primary port recovers, all traffic will again be forwarded through the primary port after a configured delay (wait-to-restore delay).

Under the Administration > Smart Pair menus you can configure the Smart Pair ports and set the wait to restore delay for a globally configured Smart Pair. Additionally you can show the Smart Pairs configured on the switch and in the show menu restore traffic manually to a configured Smart Pair.

Usage Guidelines

- ◆ Spanning-Tree must be disabled on the port in order to configure it as part of a Smart Pair.
- ◆ A configured backup port can only be a member of a single Smart Pair.
- ◆ Dynamic trunk ports cannot be configured as a Smart Pair port (only static trunk ports).
- ◆ When a trunk group is identified as a backup port of a Smart Pair, any members removed from the trunk group will not retain the Smart Pair settings.
- ◆ When all members of trunk group set as a backup port are removed, the backup port is also removed.
- ◆ The VLAN mode of the backup port should not be changed after it is configured otherwise the behavior becomes unpredictable.

Configuring the Smart Pair Global Settings

Use the Administration > Smart Pair (Configure Global) page to create a Smart Pair ID. The Smart pair ID will be used to specify two ports that are the primary and secondary members of the Smart Pair.

Parameters

These parameters are displayed:

- ◆ **Smart Pair ID** – Specifies a Smart Pair on the switch. (Default: None, Range: 1-1000 IDs can be specified.)

Web Interface

To configure global settings for LBD:

1. Click Administration, Smart Pair, Configure Global.
2. Select Add from the Action menu.
3. Enter an ID in the Smart Pair ID field.
4. Click Apply.

Figure 370: Configuring the Smart Pair Global Settings (Adding a Smart Pair)



The screenshot shows a web interface for configuring Smart Pairs. At the top, there is a breadcrumb trail: "Administration > Smart Pair". Below this, there is a "Step:" dropdown menu set to "1. Configure Global" and an "Action:" dropdown menu set to "Add". Underneath, there is a label "Smart Pair ID (1-1000)" followed by a text input field. At the bottom right, there are two buttons: "Apply" and "Revert".

Configuring Smart Pair Interface Settings

Use the Administration > Smart Pair (Configure Smart Pair Global) to add the port members of a Smart Pair. The ports must have spanning tree turned off to be available for selection.

Parameters

These parameters are displayed:

- ◆ **ID** – Selects the Smart Pair ID for configuration of its ports and wait-to-restore delay timer.
- ◆ **Primary Port** – Displays a list of Ethernet ports for selection as the primary port of a Smart Pair.
 - **Port** – Ethernet port Identifier. (Range: 1-28/54)
- ◆ **Backup Port** – Displays a list of Ethernet ports for selection as the backup port of a Smart Pair.
 - **Port** – Ethernet port Identifier. (Range: 1-28/54)
- ◆ **WTR Delay** Sets the wait-to-restore delay for a Smart Pair in seconds (Default: 30 seconds, Range: 0, 5-3600)

Web Interface

To configure the interface settings for a Smart Pair:

1. Click Administration, Smart Pair, Configure Smart Pair.
2. Select Configure from the Action menu.
3. Select the ID of the Smart Pair to be configured from the ID pull down-menu.
4. Select the Smart Pair primary port from the Primary Port pull-down menu and check the box in front of the port ID.
5. Select the Smart Pair Backup Port from the Primary Port pull-down menu and check the box in front of the port ID.
6. Input the WTR delay time in seconds and check the box in front of the field.
7. Click Apply.

Figure 371: Configuring Interfaces for a Smart Pair

Administration > Smart Pair

Steps: 2. Configure Smart Pair Action: Configure

ID: 1

Primary Port: 1

Backup Port: 1

WTR Delay: ☒ 30 seconds

Apply Revert

Show the Configured Smart Pair IDs Use the Administration > Smart Pair (Configure Global) to show the configured Smart Pair IDs.

Web Interface

To show the configured Smart Pair IDs.

1. Click Administration, Smart Pair, Configure Global.
2. Select Show from the Action menu.

Figure 372: Displaying the Smart Pair IDs.

Administration > Smart Pair

Steps: 1. Configure Global Action: Show

Smart Pair ID List: Total 1

ID
1

Delete Revert

Display the Configured Smart Pair Port Members and Restore the Traffic Use the Administration > Smart Pair (Configure Smart Pair Global) to display the port members of a Smart Pair.

Web Interface

To configure the interface settings for a Smart Pair:

1. Click Administration, Smart Pair, Configure Smart Pair.
2. Select Configure from the Show menu.
3. Select the Smart Pair ID from the ID pull-down menu.
4. Click the Restore button to manually restore traffic to the primary port of a specified Smart Pair.

Figure 373: Displaying Port Members of a Smart Pair



Multicast Filtering


This chapter describes how to configure the following multicast services:

- ◆ [IGMP Snooping](#) – Configures snooping and query parameters.
- ◆ [IGMP 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.
- ◆ [MLD Filtering and Throttling](#) – Filters specified multicast service, or throttles the maximum of multicast groups allowed on an interface.
- ◆ [Filtering MLD Query Packets on an Interface](#) – Configures the interface to drop MLD query packets.
- ◆ [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 374: 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).

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 578](#)) 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 582](#)). 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. 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 578](#)).

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 582](#)). 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 584](#)).

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.

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: Disabled)

When IGMP snooping is enabled globally, the per VLAN interface settings for IGMP snooping take precedence (see [“Setting IGMP Snooping Status per Interface” on page 586](#)).

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.

- ◆ **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)
- ◆ **Router Port Mode** – Configures multicast router ports to forward multicast streams only when multicast groups are joined. (Options: Forward, Dynamic; Default: Forward)
- ◆ **Report Suppression** – Enables report suppression on VLANs, which intercepts, absorbs and summarizes IGMP reports coming from downstream hosts. All multicast group membership states are included in a single proxy report, thereby limiting the amount of IGMP traffic sent to neighboring multicast routers. (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 375: Configuring General Settings for IGMP Snooping

Multicast > IGMP Snooping > General

IGMP Snooping Status	☐ Enabled
Proxy Reporting Status	☐ Enabled
TCR Flood	☐ Enabled
TCR Query Select	☐ Enabled
Router Alert Option	☐ Enabled
Unregistered Data Flooding	☐ Enabled
Version Exclusive	☐ Enabled
IGMP Unsolicited Report Interval (1-65535)	400 seconds
Router Port Expire Time (1-65535)	300 seconds
IGMP Snooping Version (1-2)	2
Querier Status	☐ Enabled
Router Port Mode	Forward
Report Suppression	☐ 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 578](#)) 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 376: Configuring a Static Interface for a Multicast Router

The screenshot shows a web interface for configuring a static interface for a multicast router. The breadcrumb navigation at the top reads "Multicast > IGMP Snooping > Multicast Router". Below this, there is a section titled "Action:" with a dropdown menu set to "Add Static Multicast Router". Underneath, there is a "VLAN" dropdown menu set to "1". Below that, there is an "Interface" section with two radio buttons: "Port" (selected) and "Trunk". The "Port" radio button is selected, and next to it is a dropdown menu set to "1". 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 377: Showing Static Interfaces Attached a Multicast Router




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 378: Showing Current Interfaces Attached a Multicast Router



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 578](#)). 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 379: 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 titled "Action:" with a dropdown menu set to "Add Static Member". Underneath, there are three rows of configuration fields: "VLAN" with a dropdown menu showing "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 380: Showing Static Interfaces Assigned to a Multicast Service



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 578](#).

Command Usage

Multicast Router Discovery

There have been many mechanisms used in the past to identify multicast routers. This has lead to interoperability issues between multicast routers and snooping switches from different vendors. In response to this problem, the Multicast Router Discovery (MRD) protocol has been developed for use by IGMP snooping and multicast routing devices. MRD is used to discover which interfaces are attached to multicast routers, allowing IGMP-enabled devices to determine where to send multicast source and group membership messages. (MRD is specified in draft-ietf-magma-mrdisc-07.)

Multicast source data and group membership reports must be received by all multicast routers on a segment. Using the group membership protocol query messages to discover multicast routers is insufficient due to query suppression. MRD therefore provides a standardized way to identify multicast routers without relying on any particular multicast routing protocol.



Note: The default values recommended in the MRD draft are implemented in the switch.

Multicast Router Discovery uses the following three message types to discover multicast routers:

- ◆ Multicast Router Advertisement – Advertisements are sent by routers to advertise that IP multicast forwarding is enabled. These messages are sent unsolicited periodically on all router interfaces on which multicast forwarding is enabled. They are sent upon the occurrence of these events:
 - Upon the expiration of a periodic (randomized) timer.
 - As a part of a router's start up procedure.
 - During the restart of a multicast forwarding interface.
 - On receipt of a Solicitation message.
- ◆ Multicast Router Solicitation – Devices send Solicitation messages in order to solicit Advertisement messages from multicast routers. These messages are used to discover multicast routers on a directly attached link. Solicitation messages are also sent whenever a multicast forwarding interface is initialized or re-initialized. Upon receiving a solicitation on an interface with IP multicast forwarding and MRD enabled, a router will respond with an Advertisement.
- ◆ Multicast Router Termination – These messages are sent when a router stops IP multicast routing functions on an interface. Termination messages are sent by multicast routers when:
 - Multicast forwarding is disabled on an interface.
 - An interface is administratively disabled.
 - The router is gracefully shut down.

Advertisement and Termination messages are sent to the All-Snoopers multicast address. Solicitation messages are sent to the All-Routers multicast address.



Note: MRD messages are flooded to all ports in a VLAN where IGMP snooping or routing has been enabled. To ensure that older switches which do not support MRD can also learn the multicast router port, the switch floods IGMP general query packets, which do not have a null source address (0.0.0.0), to all ports in the attached VLAN. IGMP packets with a null source address are only flooded to all ports in the VLAN if the system is operating in multicast flooding mode, such as when a new VLAN or new router port is being established, or a spanning tree topology change has occurred. Otherwise, this kind of packet is only forwarded to known multicast routing ports.

Parameters

These parameters are displayed:

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

- ◆ **IGMP Snooping Status** – When enabled, the switch will monitor network traffic on the indicated VLAN interface to determine which hosts want to receive multicast traffic. This is referred to as IGMP Snooping. (Default: Disabled)

When IGMP snooping is enabled globally (see [page 578](#)), 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 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 578](#)), or as a proxy host when IGMP snooping proxy reporting is enabled ([page 578](#)).

- ◆ **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: 100 seconds)

This attribute applies when the switch is serving as the querier ([page 578](#)), or as a proxy host when IGMP snooping proxy reporting is enabled ([page 578](#)).

- ◆ **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: 10 seconds)

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 578](#)) or IGMP querier is enabled ([page 578](#)).

- ◆ **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).

- ◆ **Report Suppression** – Enables report suppression on the VLAN, which intercepts, absorbs and summarizes IGMP reports coming from downstream hosts. All multicast group membership states are included in a single proxy report, thereby limiting the amount of IGMP traffic sent to neighboring multicast routers. (Default: Disabled)

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 381: 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

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 (10-31744): 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: 8.8.8.8

Report Suppression: ☒ Enabled

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 382: Showing Interface Settings for IGMP Snooping

VLAN Snooping	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	Report Suppression
1	Disabled	Disabled	125	180	10	2	0.0.0.0	Using global status (Disabled)	Disabled	Disabled	Using global status (Disabled)	Using global version (2)	Disabled

Filtering IGMP Packets on an Interface

Use the Multicast > IGMP Snooping > Interface (Configure Interface) page to configure an interface to drop IGMP query packets or multicast data packets, or enable IGMP authentication.

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).
- ◆ **IGMP Authentication** – Enables IGMP authentication on the specified interface. When enabled and an IGMP JOIN request is received, an authentication request is sent to a configured RADIUS server.

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 383: Dropping IGMP Query or Multicast Data Packets

Port	IGMP Query Drop	Multicast Data Drop	IGMP Authentication
1	☐ Enabled	☐ Enabled	☐ Enabled
2	☐ Enabled	☐ Enabled	☐ Enabled
3	☐ Enabled	☐ Enabled	☐ Enabled
4	☐ Enabled	☐ Enabled	☐ Enabled
5	☐ Enabled	☐ 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 578](#)).

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.

Figure 384: Showing Multicast Groups Learned by IGMP Snooping



Multicast > IGMP Snooping > Forwarding Entry

IGMP Snooping Forwarding Entry List Total: 10

VLAN	Group Address	Interface	Up Time	Expire	Count
1	224.1.1.1	Eth 1/8 (Router Port)	00:00:00:46		2 (Port)
		Eth 1/11 (Member Port)	00:00:00:46	03:46	1 (Host)
1	224.1.1.2	Eth 1/8 (Router Port)		02:24	1 (Port)
2	224.1.1.3	Eth 1/8 (Router Port)	00:00:00:14		1 (Port)
		Eth 1/9 (Router Port)			
2	238.255.255.255	Eth 1/9 (Router Port)	00:00:00:47		2 (Port)
		Eth 1/11 (Member Port)	00:00:00:47	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. There are four Action pages with separate statistics:

- Show Query Statistics
- Show VLAN Statistics
- Show Port Statistics
- Show Trunk Statistics

Statistic Pages

These parameters are displayed:

Show VLAN Statistics

- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Port** – Port identifier. (Range: 1-28/54)
- ◆ **Trunk** – Trunk identifier. (Range: 1-28)

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 385: Displaying IGMP Snooping Statistics – Query

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 386: Displaying IGMP Snooping Statistics – VLAN

The screenshot shows the 'Multicast > IGMP Snooping > Statistics' page. The 'Action' dropdown is set to 'Show VLAN Statistics'. A 'VLAN' dropdown is set to '1'. The page displays two sections: 'Input Statistics' and 'Output Statistics'. Each section has a table with four rows: 'Report', 'Leave', 'G Query', and 'G(-S)-S Query'. The 'Drop' column is only present in the 'Input Statistics' section. All values are 0. At the bottom right are 'Clear' and 'Refresh' buttons.

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

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 387: Displaying IGMP Snooping Statistics – Port

The screenshot shows the 'Multicast > IGMP Snooping > Statistics' page. The 'Action' dropdown is set to 'Show Port Statistics'. A 'Port' dropdown is set to '1'. The page displays two sections: 'Input Statistics' and 'Output Statistics'. Each section has a table with four rows: 'Report', 'Leave', 'G Query', and 'G(-S)-S Query'. The 'Drop' column is only present in the 'Input Statistics' section. All values are 0. At the bottom right is a 'Refresh' button.

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

To display IGMP snooping protocol-related statistics for a trunk:

1. Click Multicast, IGMP Snooping, Statistics.
2. Select Show Trunk Statistics from the Action list.
3. Select a Trunk.

Figure 388: Displaying IGMP Snooping Statistics – Trunk

The screenshot shows a web interface for displaying IGMP Snooping statistics for a specific trunk. The breadcrumb trail is "Multicast > IGMP Snooping > Statistics". Below this, there is an "Action:" dropdown menu set to "Show Trunk Statistics". A "Trunk:" dropdown menu is set to "1". The statistics are organized into two sections: "Input Statistics" and "Output Statistics". Each section contains a table of statistics.

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	Drop	0
Leave	0	Group	0
G Query	0		
G S S Query	0		

At the bottom right of the statistics area, there are two buttons: "Clear" and "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 389: Enabling IGMP Filtering and Throttling



Configuring IGMP Filter Profiles

Use the Multicast > IGMP Snooping > Filter (Configure Profile – Add) page to create an IGMP profile and set its access mode. Then use the (Add Multicast Group Range) page to configure the multicast groups to filter.

Command Usage

Specify a range of multicast groups by entering a start and end IP address; or specify a single multicast group by entering the same IP address for the start and end of the range.

Parameters

These parameters are displayed:

Add

- ◆ **Profile ID** – Creates an IGMP profile. (Range: 1-4294967295)
- ◆ **Access Mode** – Sets the access mode of the profile; either permit or deny. (Default: Deny)

When the access mode is set to permit, IGMP join reports are processed when a multicast group falls within the controlled range. When the access mode is set to deny, IGMP join reports are only processed when the multicast group is not in the controlled range.

Add Multicast Group Range

- ◆ **Profile ID** – Selects an IGMP profile to configure.
- ◆ **Start Multicast IP Address** – Specifies the starting address of a range of multicast groups.
- ◆ **End Multicast IP Address** – Specifies the ending address of a range of multicast groups.

Web Interface

To create an IGMP filter profile and set its access mode:

1. Click Multicast, IGMP Snooping, Filter.
2. Select Configure Profile from the Step list.
3. Select Add from the Action list.
4. Enter the number for a profile, and set its access mode.
5. Click Apply.

Figure 390: Creating an IGMP Filtering Profile

The screenshot shows a web interface for configuring IGMP filtering. The breadcrumb trail at the top reads "Multicast > IGMP Snooping > Filter". Below this, there are two dropdown menus: "Step:" set to "2. Configure Profile" and "Action:" set to "Add". The main configuration area contains two fields: "Profile ID (1-4294967295)" with the value "19" entered, and "Access Mode" set to "Permit" via a dropdown menu. At the bottom right of the form are two buttons: "Apply" and "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 391: Showing the IGMP Filtering Profiles Created



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 392: Adding Multicast Groups to an IGMP Filtering Profile



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 393: Showing the Groups Assigned to an IGMP Filtering Profile



Configuring IGMP Filtering and Throttling for Interfaces

Use the Multicast > IGMP Snooping > Filter (Configure Interface) page to assign an 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 394: Configuring IGMP Filtering and Throttling Interface Settings



Port	Profile ID	Max Multicast Groups (1-1024)	Current Multicast Groups	Throttling Action Mode	Throttling Status
1	1	1024	0	Deny	False
2	1	1024	0	Deny	False
3	1	1024	0	Deny	False
4	1	1024	0	Deny	False
5	1	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: 1)
- ◆ **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.)
- ◆ **Proxy Reporting** – Enables MLD Snooping with Proxy Reporting. (Default: Disabled)

When proxy reporting is enabled, reports received from downstream hosts are summarized and used to build internal membership states. Proxy-reporting devices may use the all-zeros IP source address when forwarding any summarized reports upstream. For this reason, IGMP membership reports received by the snooping switch must not be rejected because the source IP address is set to 0.0.0.0.
- ◆ **Unsolicited Report Interval** – Specifies how often the upstream interface should transmit unsolicited MLD snooping reports when proxy reporting is enabled. (Default: 400 seconds)

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 395: Configuring General Settings for MLD Snooping

Multicast > MLD Snooping > General

MLD Snooping Status	☒ Enabled
Querier Status	☒ Enabled
Robustness (2-10)	2
Query Interval (60-125)	125 seconds
Query Max Response Time (5-25)	10 seconds
Router Port Expiry Time (300-500)	300 seconds
MLD Snooping Version (1-2)	2
Unknown Multicast Mode	To Router Port ▼
Proxy Reporting	☒ Enabled
Unsolicited Report Interval (1-65535)	400 seconds

Apply Reset

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** – 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.
- ◆ **Immediate Leave By Host** – Specifies that the member port will be deleted only when there are no hosts joining this group.

If the “By Host” option is used, the router/querier will not send out a group-specific query when an MLD 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.

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 396: Configuring Immediate Leave for MLD Snooping

The screenshot shows a web interface for configuring MLD Snooping. The breadcrumb trail is 'Multicast > MLD Snooping > Interface'. Below this, there is a 'VLAN' dropdown menu. Underneath, there are two rows of settings: 'Immediate Leave' with a checked checkbox and the text 'Enabled', and 'Immediate Leave By Host' with a checked checkbox and the text 'Enabled'. At the bottom right of the form are two buttons: 'Apply' and '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 605](#)) 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 397: Configuring a Static Interface for an IPv6 Multicast Router

The screenshot shows the 'Multicast > MLD Snooping > Multicast Router' configuration page. The 'Action:' dropdown is set to 'Add Static Multicast Router'. Below this, there is a 'VLAN' dropdown menu with '1' selected. Under the 'Interface:' section, the 'Port' radio button is selected with '1' in the adjacent input field, and the 'Trunk' radio button is unselected. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 398: Showing Static Interfaces Attached an IPv6 Multicast Router


The screenshot shows the 'Multicast > MLD Snooping > Multicast Router' configuration page. The 'Action:' dropdown is set to 'Show Static Multicast Router'. Below this, there is a 'VLAN' dropdown menu with '1' selected. A table titled 'Static Multicast Router Interface List' is displayed, showing two entries. At the bottom right, there are 'Delete' and 'Revert' buttons.

	Interface
☐	Unit 1 / Port 1
☐	Unit 1 / Port 2

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 399: Showing Current Interfaces Attached an IPv6 Multicast Router



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 605](#)). 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 400: Assigning an Interface to an IPv6 Multicast Service

The screenshot shows the 'Multicast > MLD Snooping > MLD Member' configuration page. The 'Action' dropdown is set to 'Add Static Member'. The 'VLAN' is set to '1'. The 'Multicast IPv6 Address' is 'FF02:0:0:0:0:0:0:10C'. The 'Interface' is set to 'Port 1' under the 'Port' tab. There are 'Apply' and 'Reset' buttons at the bottom right.

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 401: Showing Static Interfaces Assigned to an IPv6 Multicast Service

The screenshot shows the 'Multicast > MLD Snooping > MLD Member' configuration page with the 'Action' dropdown set to 'Show Static Member'. The 'VLAN' is set to '1'. Below the configuration fields is a table titled 'MLD Member Interface List' with the following data:

	Multicast IPv6 Address	Interface
☐	FF02:0:0:0:0:0:0:101	Unit 1 / Port 1
☐	FF02:0:0:0:0:0:0:102	Unit 1 / Port 2
☐	FF01:0	Unit 1 / Port 12
☐	FF01:1	Unit 1 / Port 13
☐	FF01:2	Unit 1 / Port 14
☐	FF01:3	Unit 1 / Port 15
☐	FF01:4	Unit 1 / Port 16
☐	FF02:0:0:0:0:0:0:FF	Trunk 3

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

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 402: Showing Current Interfaces Assigned to an IPv6 Multicast Service



Multicast IPv6 Address	Interface	Type
FF02::01:01:01:01	Unit 1 / Port 1	User
FF02::01:01:01:02	Unit 1 / Port 2	User
FF01::1	Unit 1 / Port 12	User
FF01::2	Unit 1 / Port 13	Multicast Data
FF01::3	Unit 1 / Port 14	User
FF01::4	Unit 1 / Port 15	User
FF01::5	Unit 1 / Port 16	User
FF02::01:01:01:FF	Trunk 3	User
FF02::01:01:01:FF	Trunk 5	MLD Snooping

Filtering MLD Query Packets on an Interface

Use the Multicast > MLD Snooping > Query Drop page to configure an interface to drop MLDF query packets.

Parameters

These parameters are displayed:

- ◆ **Interface** – Port or trunk identifier.
- ◆ **Query Drop** – Drops any received MLD query packets. (Default: Disabled)

This feature can be used to drop any query packets received on the specified interface. If this switch is acting as a Querier, this prevents it from being affected by messages received from another Querier.


Web Interface

To drop IGMP query packets:

1. Click Multicast, MLD Snooping, Query Drop.
2. Select Port or Trunk interface.
3. Enable query drop for any interface.

4. Click Apply.

Figure 403: Dropping MLD Query Packets



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 404: Showing IPv6 Multicast Services and Corresponding Sources



Displaying MLD Snooping Statistics Use the Multicast > IGMP Snooping > Statistics pages to display MLD snooping protocol-related statistics.

Parameters

These parameters are displayed:

Input

- ◆ **Interface** – The unit/port or VLAN interface.
- ◆ **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 MLD group report received.

- ◆ **Join Success** – The number of times a multicast group was successfully joined.
- ◆ **Group** – The number of MLD groups active on this interface.

Output

Same as input parameters listed above, except that the direction of transmission is outbound.

Query

- ◆ **Other Querier Address** – 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 Address** – IP address of local querier on this interface.
- ◆ **Self Querier Expire Time** – 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 group specific queries received on this interface.
- ◆ **Specific Query Sent** – The number of group specific queries sent from this interface.

Summary

Physical Interface (Port/Trunk)

- ◆ **Number of Groups** – Number of active MLD groups active on the specified interface.
- ◆ *Querier*
 - *Transmit*
 - **General** – The number of general queries sent from this interface.

- **Group Specific** – The number of group specific queries sent from this interface.
- *Received*
 - **General** – The number of general queries received on this interface.
 - **Group Specific** – The number of group specific queries received on this interface.
- ◆ *Report & Leave*
 - *Transmit*
 - **Report** – The number of MLD membership reports sent from this interface.
 - **Leave** – The number of leave messages sent from this interface.
 - *Received*
 - **Report** – The number of MLD membership reports received on this interface.
 - **Leave** – The number of leave messages received on this interface.
 - **Join Success** – The number of times a multicast group was successfully joined.
 - **Filter Drop** – The number of messages dropped due to filter settings.
 - **Source Port Drop** – The number of dropped messages that are received on MVR source port or mrouter port.
 - **Others Drop** – The number of received invalid messages.

Logical Interface (VLAN) – The following parameters are included for a VLAN.

- ◆ **Number of Groups** – Number of active MLD groups active on the specified interface.
- ◆ *Querier*
 - **Other Querier** – IPv6 address of remote querier on this interface.
 - **Other Uptime** – Time remote querier has been up.
 - **Other Expire** – Time after which remote querier is assumed to have expired.

- **Self Addr** – IPv6 address of local querier on this interface.
- **Self Expire** – Time after which local querier is assumed to have expired.
- **Self Uptime** – Time local querier has been up.
- *Transmit*
 - **General** – The number of general queries sent from this interface.
 - **Group Specific** – The number of group specific queries sent from this interface.
- *Received*
 - **General** – The number of general queries received on this interface.
 - **Group Specific** – The number of group specific queries received on this interface.
- ◆ *Report & Leave*
 - **Host Addr** – The link-local or global IPv6 address that is assigned on that VLAN.
 - **Unsolicit Expire** – The number of group leaves resulting from timeouts instead of explicit leave messages.
 - *Transmit*
 - **Report** – The number of MLD membership reports sent from this interface.
 - **Leave** – The number of leave messages sent from this interface.
 - *Received*
 - **Report** – The number of MLD membership reports received on this interface.
 - **Leave** – The number of leave messages received on this interface.

Clear

Parameters

These parameters are displayed:


- ◆ **All** – Clears statistics for all MLD messages.
- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Unit** – Stack unit. (Range: 1)
- ◆ **Port** – Port identifier. (Range: 1-28/54)
- ◆ **Trunk** – Trunk identifier. (Range: 1-28)

Web Interface

To display MLD snooping input-related message statistics:

1. Click Multicast, MLD Snooping, Statistics.
2. Select Input.

Figure 405: Displaying MLD Snooping Statistics – Input



The screenshot shows the 'Multicast > MLD Snooping > Statistics' page. At the top, there are radio buttons for 'Input' (selected), 'Output', 'Query', 'Summary', and 'Clear'. Below this, the title 'Input Statistics' is followed by 'Total: 19'. A table displays statistics for interfaces Eth 1/1 through Eth 1/6. The table has columns for Interface, Report, Leave, G-Query, G(S)-S-Query, Drop, Join Success, and Group. All values in the table are 0.

Interface	Report	Leave	G-Query	G(S)-S-Query	Drop	Join Success	Group
Eth 1/1	0	0	0	0	0	0	0
Eth 1/2	0	0	0	0	0	0	0
Eth 1/3	0	0	0	0	0	0	0
Eth 1/4	0	0	0	0	0	0	0
Eth 1/5	0	0	0	0	0	0	0
Eth 1/6	0	0	0	0	0	0	0

To display MLD snooping output-related message statistics:

1. Click Multicast, MLD Snooping, Statistics.
2. Select Output.

Figure 406: Displaying MLD Snooping Statistics – Output

Multicast > MLD Snooping > Statistics

Type: ☐ Input ☒ Output ☐ Query ☐ Summary ☐ Clear

Output Statistics: Total: 10

Interface	Request	Leave	IG Query	IG-S Query	Drop	Group
Eth 1/1	0	0	0	0	0	0
Eth 1/2	0	0	0	0	0	0
Eth 1/3	0	0	0	0	0	0
Eth 1/4	0	0	0	0	0	0
Eth 1/5	0	0	0	0	0	0
Eth 1/6	0	0	0	0	0	0

To display MLD query message statistics:

1. Click Multicast, MLD Snooping, Statistics.
2. Select Query.

Figure 407: Displaying MLD Snooping Statistics – Query

Multicast > MLD Snooping > Statistics

Type: ☐ Input ☐ Output ☒ Query ☐ Summary ☐ Clear

VLAN: 1 ▼

Query Statistics

Other Querier Address	None
Other Querier Expire	0(m):0(s)
Other Querier Uptime	0(h):0(m):0(s)
Self Querier Address	::
Self Querier Expire Time	0(m):0(s)
Self Querier Uptime	0(h):0(m):0(s)
General Query Received	0
General Query Sent	0
Specific Query Received	0
Specific Query Sent	0

To display MLD summary statistics for a port or trunk:

1. Click Multicast, MLD Snooping, Statistics.
2. Select Summary.
3. Select a port or trunk.

Figure 408: Displaying MLD Snooping Statistics – Summary (Port/Trunk)

Multicast > MLD Snooping > Statistics

Type ☐ Input ☐ Output ☐ Query ☒ Summary ☐ Clear

Interface ☐ VLAN ☐ Unit / Port ☐ Trunk

Summary Statistics

Number of Groups	0		
Queried		Report & Leave	
Transmit		Transmit	
General	0	Report	0
Group Specific	0	Leave	0
Received		Received	
General	0	Report	0
Group Specific	0	Leave	0
		Join Success	0
		Source Port Drop	0
		Others Drop	0

To display MLD summary statistics for a VLAN:

1. Click Multicast, MLD Snooping, Statistics.
2. Select Summary.
3. Select a VLAN.

Figure 409: Displaying MLD Snooping Statistics – Summary (VLAN)

Multicast > MLD Snooping > Statistics

Type ☐ Input ☐ Output ☐ Query ☒ Summary ☐ Clear

Interface ☒ VLAN

☐ Unit / Port

☐ Trunk

Summary Statistics

Number of Groups		0	
Querier			
Other Querier	None	Report & Leave	
Other Uptime	0(h):0(m):0(s)	Host Addr	None
Other Expire	0(m):0(s)	Unsolicit Expire	0 sec
Self Addr	None		
Self Expire	0(m):0(s)		
Self Uptime	0(h):0(m):0(s)		
Transmit		Transmit	
General	0	Report	0
Group Specific	0	Leave	0
Received		Received	
General	0	Report	0
Group Specific	0	Leave	0
		Join Success	0
		Source Port Drop	0
		Others Drop	0

To clear MLD statistics:

1. Click Multicast, MLD Snooping, Statistics.
2. Select Clear.
3. Select All or enter the required interface.
4. Click Clear.

Figure 410: Clearing MLD Snooping Statistics



Filtering and Throttling MLD 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 MLD filtering feature fulfills this requirement by restricting access to specified multicast services on a switch port, and MLD throttling limits the number of simultaneous multicast groups a port can join.

MLD filtering enables you to assign a profile to a switch port that specifies multicast groups that are permitted or denied on the port. An MLD 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, MLD join reports received on the port are checked against the filter profile. If a requested multicast group is permitted, the MLD join report is forwarded as normal. If a requested multicast group is denied, the MLD join report is dropped.

MLD 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 MLD 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 MLD Filtering and Throttling Use the Multicast > MLD Snooping > Filter (Configure General) page to enable IGMP filtering and throttling globally on the switch.

Parameters

These parameters are displayed:

- ◆ **MLD Filter Status** – Enables MLD filtering and throttling globally for the switch. (Default: Disabled)

Web Interface

To enable MLD filtering and throttling on the switch:

1. Click Multicast, MLD Snooping, Filter.
2. Select Configure General from the Step list.
3. Enable MLD Filter Status.
4. Click Apply.

Figure 411: Enabling MLD Filtering and Throttling



Configuring MLD Filter Profiles Use the Multicast > MLD Snooping > Filter (Configure Profile – Add) page to create an MLD 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 IPv6 address; or specify a single multicast group by entering the same IPv6 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, MLD join reports are processed when a multicast group falls within the controlled range. When the access mode is set to deny, MLD 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 IPv6 Address** – Specifies the starting address of a range of multicast groups.
- ◆ **End Multicast IPv6 Address** – Specifies the ending address of a range of multicast groups.

Web Interface

To create an MLD filter profile and set its access mode:

1. Click Multicast, MLD 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 412: Creating an MLD Filtering Profile

The screenshot shows a web interface for configuring MLD filtering. The breadcrumb navigation at the top reads "Multicast > MLD Snooping > Filter". Below this, there are two dropdown menus: "Step:" with "2. Configure Profile" selected, and "Action:" with "Add" selected. The main configuration area contains two fields: "Profile ID (1-4294967295)" with the value "19" entered, and "Access Mode" with a dropdown menu showing "Permit". At the bottom right of the form are two buttons: "Apply" and "Revert".

To show the MLD filter profiles:

1. Click Multicast, MLD Snooping, Filter.
2. Select Configure Profile from the Step list.
3. Select Show from the Action list.

Figure 413: Showing the MLD Filtering Profiles Created



To add a range of multicast groups to an MLD filter profile:

1. Click Multicast, MLD 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 414: Adding Multicast Groups to an MLD Filtering Profile



To show the multicast groups configured for an MLD filter profile:

1. Click Multicast, MLD 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 415: Showing the Groups Assigned to an MLD Filtering Profile

Multicast > MLD Snooping > Filter

Step: 2. Configure Profile Action: Show Multicast Group Range

Profile ID: 15

Multicast IPv6 Address Range List Total: 1

Start Multicast IPv6 Address	End Multicast IPv6 Address
ff01::101	ff01::202

Delete Revert

Configuring MLD Filtering and Throttling for Interfaces

Use the Multicast > MLD Snooping > Filter (Configure Interface) page to assign an MLD 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

- ◆ MLD 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 MLD 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 MLD 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: 1-255; Default: 255)

- ◆ **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 MLD filtering or throttling for a port or trunk:

1. Click Multicast, MLD 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 416: Configuring MLD Filtering and Throttling Interface Settings



The screenshot shows the 'Multicast > MLD Snooping > Filter' web interface. The 'Step' dropdown is set to '1. Configure Interface'. The 'Interface' section has radio buttons for 'Port' (selected) and 'Trunk'. Below this is a table titled 'MLD Filter and Throttling Port List Total: 28'. The table has six columns: Port, Profile ID, Max Multicast Groups (1-255), Current Multicast Groups, Throttling Action Mode, and Throttling Status. There are five rows of data, all with 'Deny' as the action mode and 'False' as the throttling status.


Port	Profile ID	Max Multicast Groups (1-255)	Current Multicast Groups	Throttling Action Mode	Throttling Status
1	▼	255	0	Deny ▼	False
2	▼	255	0	Deny ▼	False
3	▼	255	0	Deny ▼	False
4	▼	255	0	Deny ▼	False
5	▼	255	0	Deny ▼	False

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 the 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 417: 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 631](#)).
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 633](#)).
3. Set the interfaces that will join the MVR as source ports or receiver ports (see [“Configuring MVR Interface Status” on page 635](#)).
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 638](#)).

- ◆ 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 418: Configuring Global Settings for MVR

The screenshot shows the 'Multicast > MVR' configuration page. At the top, there is a 'Step' dropdown menu set to '1. Configure Global'. Below this, there are four configuration fields: 'Proxy Switching' with a checked 'Enabled' checkbox, 'Robustness Value (1-255)' with a text input field containing '1', 'Proxy Query Interval (2-31744)' with a text input field containing '125' and a 'sec' unit label, and 'Source Port Mode' with a dropdown menu set to 'Always Forward'. At the bottom right of the form are 'Apply' and 'Revert' buttons.

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 167](#)), 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 635](#)).

◆ **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 419: Configuring Domain Settings for MVR

The screenshot shows a web interface for configuring MVR settings. The title bar reads 'Multicast > MVR'. Below it, a 'Step:' dropdown menu is set to '2. Configure Domain'. The configuration area contains several fields: 'Domain ID' with a dropdown showing '1', 'MVR Status' with a checkbox labeled 'Enabled' that is checked, 'MVR VLAN' with a dropdown showing '1', 'MVR Running Status' with the text 'Inactive', 'MVR Current Learned Groups' with the value '0', 'Forwarding Priority (0-7)' with a slider set to '0', and 'Upstream Source IP' with a text box containing '0.0.0.0'. At the bottom right, there are 'Apply' and 'Reset' buttons.

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 420: Configuring an MVR Group Address Profile

Multicast > MVR

Step: 2: Configure Profile Action: Add

Profile Name: sales

Start IP Address: 234.5.6.8

End IP Address: 234.5.6.10

Apply Revert

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 421: Displaying MVR Group Address Profiles

Multicast > MVR

Step: 3: Configure Profile Action: Show

MVR Profile List Total: 2

	Profile Name	Start IP Address - End IP Address
☐	profile1	224.1.1.1 - 224.1.1.19
☐	profile2	224.1.2.1 - 224.1.2.100

Done Revert

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 422: 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 423: 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
all	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 638](#)).

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 167](#)).

- ◆ 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 638](#)).

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 167](#)).
 - **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 638](#)).
 - **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)
 - **By Group** – The receiver port is immediately removed from the multicast group identified in the leave message.
 - **By Host IP** – The router/querier will not send out a group-specific query when an IGMPv2/v3 leave message is received (the same as it would without this option having been used). Instead of immediately deleting that group, it will look up the record, and only delete the group if there are no other subscribers for it on the member port. Only when all hosts on that port leave the group will the member port be deleted.

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 424: Configuring Interface Settings for MVR

Step: 3. Configure Interface

Domain ID: 1

Interface: ☒ Port ☐ Trunk

Port List Total: 25

Port	Type	Forwarding Status	MVR Status	Immediate Leave
1	Non-MVR	Forwarding	Inactive	☐ By Group
2	Non-MVR	Discarding	Inactive	☐ By Group
3	Non-MVR	Discarding	Inactive	☐ By Group
4	Non-MVR	Discarding	Inactive	☐ By Group
5	Non-MVR	Discarding	Inactive	☐ By Group

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 425: 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	224.5.6.8

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 426: Showing the Static MVR Groups Assigned to a Port



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** – Indicates the source address of the multicast service (these entries are marked as "Source"), or displays an asterisk if the group address has been statically assigned. Also shows the interfaces with subscribers for multicast services provided through the MVR 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 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 427: Displaying MVR Receiver Groups



Group IP Address	VLAN	Port	Up Time	Expire	Count
224.1.1.1	2	Unit 1 / Port 1 (Source)	00:00:30		2 (Port)
	1	Unit 1 / Port 2 (Receiver)	00:01:10	00:00	4 (Host)
224.1.1.2	4	Unit 1 / Port 3 (Source)	00:00:50		4 (Port)
	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. There are four Action pages with separate statistics:

- Show Query Statistics
- Show VLAN Statistics
- Show Port Statistics
- Show Trunk Statistics

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Port** – Port identifier. (Range: 1-28/54)
- ◆ **Trunk** – Trunk identifier. (Range: 1-28)

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 428: Displaying MVR Statistics – Query

The screenshot shows 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 is a 'Domain ID' dropdown menu 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 with a tooltip 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, Port or Trunk:

1. Click Multicast, MVR.
2. Select Show Statistics from the Step list.
3. Select Show VLAN/Port/Trunk Statistics from the Action list.
4. Select an MVR domain.
5. Select a VLAN/Port/Trunk.

Figure 429: Displaying MVR Statistics – VLAN

The screenshot shows a web interface for Multicast > MVR configuration. At the top, there is a breadcrumb 'Multicast > MVR'. Below it, there are two dropdown menus: 'Step: 7. Show Statistics' and 'Action: Show VLAN Statistics'. Under these, there are two more dropdown menus: 'Domain ID: 1' and 'VLAN: 1'. The main content area is divided into two sections: 'Input Statistics' and 'Output Statistics'. Each section has a table of statistics. In the 'Input Statistics' section, the values for Report, Leave, G Query, and G(-S)-S Query are all 0. In the 'Output Statistics' section, the values for Report, Leave, G Query, and G(-S)-S Query are also all 0. To the right of the 'Input Statistics' table, there are three more statistics: Drop (0), Join Success (0), and Group (0). At the bottom right of the page, there is a 'Refresh' button.

Input Statistics		Output Statistics	
Report	0	Report	0
Leave	0	Leave	0
G Query	0	G Query	0
G(-S)-S Query	0	G(-S)-S Query	0

Drop	
Drop	0
Join Success	0
Group	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 430: 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 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		Output Statistics	
Report	0	Report	0
Leave	0	Leave	0
G Query	0	G Query	0
G(-S)-S Query	0	G(-S)-S Query	0
Drop	0	Join Success	0
Join Success	0	Group	0
Group	0		

At the bottom right of the statistics section 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 628](#)).

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 648](#)).
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 649](#)).

3. Set the interfaces that will join the MVR as source ports or receiver ports (see [“Configuring MVR6 Interface Status” on page 652](#)).
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 654](#)).

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.
 - **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. (This is the default setting.)

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 431: Configuring Global Settings for MVR6

The screenshot shows the 'Multicast > MVR6' configuration page. At the top, there is a breadcrumb 'Multicast > MVR6' and a 'Step: 1. Configure Global' dropdown. Below this, there are four configuration fields: 'Proxy Switching' with a checked 'Enabled' checkbox, 'Robustness Value (1-18)' with a text input field containing '1', 'Proxy Query Interval (2-31744)' with a text input field containing '125' and a 'sec' unit label, and 'Source Port Mode' with a dropdown menu set to 'Always Forward'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 167](#)), 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 652](#)).
- ◆ **MVR6 Current Learned Groups** – The number of MVR6 groups currently assigned to this domain.
- ◆ **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 432: Configuring Domain Settings for MVR6

The screenshot shows a web interface for configuring MVR6 domains. The breadcrumb is 'Multicast > MVR6'. The 'Step' dropdown is set to '2. Configure Domain'. The configuration fields are as follows:

Domain ID	1
MVR6 Status	☒ Enabled
MVR6 VLAN	1
MVR6 Running Status	Inactive
MVR6 Current Learned Groups	0
Upstream Source IPv6	

At the bottom right, there are 'Apply' and 'Reset' 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: 1-20 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 433: Configuring an MVR6 Group Address Profile

The screenshot shows a web interface for configuring an MVR6 group address profile. The breadcrumb is 'Multicast > MVR6'. The 'Step' dropdown is set to '2. Configure Profile' and the 'Action' dropdown is set to 'Add'. There are three input fields: 'Profile Name' with the value 'p0', 'Start IPv6 Address' with the value 'ff00::1', and 'End IPv6 Address' with the value 'ff00::9'. At the bottom right, there are two buttons: 'Apply' and '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 434: Displaying MVR6 Group Address Profiles



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 435: Assigning an MVR6 Group Address Profile to a Domain



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 436: Showing MVR6 Group Address Profiles Assigned to a Domain



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).

- ◆ 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 638).

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 171).

- ◆ 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 638).

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 167](#)).
 - **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 167](#)).
- ◆ **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)
 - **By Group** – The receiver port is immediately removed from the multicast group identified in the MLD snooping leave message.
 - **By Host IP** – The router/querier will not send out a group-specific query when an MLD snooping leave message is received (the same as it would without this option having been used). Instead of immediately deleting that group, it will look up the record, and only delete the group if there are no other subscribers for it on the member port. Only when all hosts on that port leave the group will the member port be deleted.

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 437: Configuring Interface Settings for MVR6

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 438: Assigning Static MVR6 Groups to a Port

The screenshot shows the 'Multicast > MVR6' configuration page. At the top, there is a 'Step' dropdown menu set to '5. Configure Static Group Member' and an 'Action' dropdown menu set to 'Add'. Below these, there are several input fields: 'Domain ID' with a dropdown menu showing '1', 'Interface' with radio buttons for 'Port' (selected) and 'Trunk', and a dropdown menu showing '2', 'VLAN' with a dropdown menu showing '2', and 'Group IPv6 Address' with a text input field containing 'ff02::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 439: Showing the Static MVR6 Groups Assigned to a Port



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 440: Displaying MVR6 Receiver Groups



Group IPv6 Address	VLAN	Port	Up-Time	Expires	Count
FF02::1	1	Unit 1 / Port 1 (Source)	00:00:11:35		2 (Port)
	1	Unit 1 / Port 2 (Receiver)			3 (Host)

Displaying MVR6 Statistics

Use the Multicast > MVR6 > Show Statistics pages to display MVR6 protocol-related statistics for the specified interface. There are four Action pages with separate statistics:

- Show Query Statistics
- Show VLAN Statistics
- Show Port Statistics
- Show Trunk Statistics

Parameters

These parameters are displayed:

Query Statistics

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **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 or Trunk Statistics

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Port** – Port identifier. (Range: 1-28/54)
- ◆ **Trunk** – Trunk identifier. (Range: 1-16)

Input Statistics

- **Report** – The number of MLD membership reports received on this interface.
- **Done** – The number of messages received by this multicast router for a specific multicast group, to signal the router that there may not be any further group members on the local subnet.
- **Leave** – The number of leave messages sent from 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.

Output Statistics

- **Report** – The number of MLD membership reports sent from this interface.
- **Done** – The number of messages sent by a host for a specific multicast group, to signal the router that there may not be any further group members on the local subnet.
- **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.
- **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.

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 441: Displaying MVR6 Statistics – Query

The screenshot shows a web interface for configuring Multicast VLAN Registration (MVR6). The breadcrumb navigation at the top reads "Multicast > MVR6". Below this, there are two dropdown menus: "Step:" set to "7. Show Statistics" and "Action:" set to "Show Query Statistics". A "Domain ID" dropdown is set to "1". Under the heading "Query Statistics", a table displays various metrics. The metrics and their values are: Querier IPv6 Address (empty), 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), and 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.

Query Statistics	
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

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 442: 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: 1' and 'VLAN: 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	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

At the bottom right of the statistics area, there are two buttons: 'Clear' and '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 443: Displaying MVR6 Statistics – Port

The screenshot shows a web interface for configuring Multicast VLAN Registration (MVR6). The breadcrumb trail is "Multicast > MVR6". Below this, there are two dropdown menus: "Step:" set to "7. Show Statistics" and "Action:" set to "Show Port Statistics".

Below the dropdowns, there are two more dropdown menus: "Domain ID" and "Port", both set to "1".

The main content area displays statistics for the selected domain and port. It is divided into two sections: "Input Statistics" and "Output Statistics".

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

At the bottom right of the statistics section, there are two buttons: "Clear" and "Refresh".

To display MVR6 protocol-related statistics for a trunk:

1. Click Multicast, MVR6.
2. Select Show Statistics from the Step list.
3. Select Show Trunk Statistics from the Action list.
4. Select an MVR6 domain.
5. Select a Trunk.

Figure 444: Displaying MVR6 Statistics – Trunk

The screenshot shows a web interface for Multicast > MVR6. At the top, there are two dropdown menus: 'Step: 5. Show Statistics' and 'Action: Show Trunk Statistics'. Below these, there are two more dropdown menus: 'Domain ID: 1' and 'Trunk: 1'. The main content area is divided into two sections: 'Input Statistics' and 'Output Statistics'. Each section has a table of statistics with values of 0.

Input Statistics	
Report	0
Done	0
G-Query	0
G(-S)-S Query	0
Drop	0
Join Success	0
Group	0

Output Statistics	
Report	0
Done	0
G-Query	0
G(-S)-S Query	0

At the bottom right, there are two buttons: 'Clear' and 'Refresh'.

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 proxy ARP or static addresses, and how to display entries in the ARP cache.

Using the Ping Function

Use the Tools > Ping page to send ICMP echo request packets to another node on the network.

Parameters

These parameters are displayed:

- ◆ **Host Name/IP Address** – Alias or IPv4/IPv6 address of the host.
- ◆ **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 Tools, Ping.
2. Specify the target device and ping parameters.
3. Click Apply.

Figure 445: Pinging a Network Device

The screenshot displays the 'Tool > Ping' web interface. At the top, there's a title bar 'Tool > Ping'. Below it, there are input fields for 'Host Name/IP Address', 'Probe Count (1-16)' (set to 5), and 'Data Size (IPv4: 32-512, IPv6: 8-1500)' (set to 32 bytes). A note below these fields explains the data size constraints: 'Note: For IPv4 Data Size: 0 - 31 changed to 32 bytes; 32 - 512 is valid input; 513 - 1500 changed to 512 bytes; < 0 or > 1500 not valid input'. There are 'Apply' and 'Revert' buttons. The 'Result' section shows the output of the ping command: 'PING to 192.168.2.99, by 5 of 32-byte payload ICMP packets, timeout is 3 seconds'. It lists five 'response time: 0 ms' entries. Below that, it shows 'Ping statistics for 192.168.2.99: 5 packets transmitted, 5 packets received (100%), 0 packets lost (0%)' and 'Approximate round trip times: Minimum = 0 ms, Maximum = 0 ms, Average = 0 ms'.

Using the Trace Route Function

Use the Tools > Trace Route page to show the route packets take to the specified destination.

Parameters

These parameters are displayed:

- ◆ **Host Name/IP Address** – Alias or 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 Tools, 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 667), 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 39: 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 Tools > 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 Tools, ARP.
2. Select Configure General from the Step List.
3. 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 the 'Tools > ARP' configuration page. At the top, there's a 'Step:' dropdown menu set to '1. Configure General'. Below this, there are three main configuration sections: 'Timeout (386-86400)' with a text input field containing '1200' and a unit dropdown set to 'sec'; 'Proxy ARP' with a status dropdown set to 'Enabled'; and 'VLAN' with a dropdown menu set to '1'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

Configuring Static ARP Addresses

For devices that do not respond to ARP requests or do not respond in a timely manner, traffic will be dropped because the IP address cannot be mapped to a physical address. If this occurs, use the Tools > 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-xx-xx or xxxxxxxxxxxx)

Web Interface

To map an IP address to the corresponding physical address in the ARP cache:

1. Click Tools, 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 a web interface for configuring static ARP entries. At the top, it says 'Tools > ARP'. Below that, there's a 'Step:' dropdown menu set to '2. Configure Static Address' and an 'Action:' dropdown menu set to 'Add'. The main area has two input fields: 'IP Address' with the value '10.2.78.106' and 'MAC Address' with the value '00-e0-b5-00-00-83'. To the right of the MAC field is a placeholder text: '(xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx)'. At the bottom right, there are two buttons: 'Apply' and 'Reset'.

To display static entries in the ARP cache:

1. Click Tools, ARP.
2. Select Configure Static Address from the Step List.
3. Select Show from the Action List.

Figure 450: Displaying Static ARP Entries



Displaying Dynamic or Local ARP Entries

Use the Tools > ARP 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 Tools, ARP.
2. Select Show Information from the Step List.
3. Click ARP Addresses.

Figure 451: Displaying ARP Entries



Displaying ARP Statistics Use the Tools > ARP (Show Information) page to display statistics for ARP messages crossing all interfaces on this switch.

Parameters

These parameters are displayed:

Table 40: 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 Tools, ARP.
2. Select Show Information from the Step List.
3. Click Statistics.

Figure 452: Displaying ARP Statistics



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 either be manually configured or dynamically generated.

This chapter provides information on network functions including:

- ◆ [IPv4 Configuration](#) – Sets an IPv4 address for management access.
- ◆ [IPv6 Configuration](#) – Sets an IPv6 address for management access.

Setting the Switch's IP Address (IP Version 4)

This section describes how to configure an initial IPv4 interface for management access over the network, and how to create an interface to multiple subnets. This switch supports both IPv4 and IPv6, and can be managed 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 global unicast or link-local address can be manually configured, or a link-local address can be dynamically generated. For information on configuring the switch with an IPv6 address, see [“Setting the Switch's IP Address \(IP Version 6\)” on page 677](#).

Configuring IPv4 Interface Settings

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. To configure this device as the default gateway, use the IP > Routing > Static Routes (Add) page, set the destination address to the required interface, and the next hop to null address 0.0.0.0.

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 can configure static routes ([page 702](#)), or a default gateway using the IP > Routing > Static Routes (Add) page (see [“Configuring Static Routes” on page 702](#)) or the IP > IPv6 Configuration (Configure Global) page (see [“Configuring the IPv6 Default Gateway” on page 677](#)).
- ◆ The precedence for configuring IP interfaces is the IP > General > Routing Interface (Add Address) menu, and then static routes ([page 702](#)).

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. (Default: VLAN 1)
- ◆ **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 453: Configuring a Static IPv4 Address

IP > General > Routing Interface

Action: Add Address

VLAN: 1

IP Address Mode: User Specified

IP Address Type: Primary

IP Address:

Subnet Mask:

Restart DHCP Click this button to request DHCP client request

Apply Revert

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 454: 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 resend DHCP client request.

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 455: Showing the Configured IPv4 Address for an Interface

IP > General > Routing Interface

Action: Show Address ▼

VLAN: 1 ▼

IP Address Mode: User Specified

Routing Interface IP List Total: 1

	IP Address Type	IP Address	Subnet Mask
☐	Primary	192.168.1.1	255.255.255.0

Delete 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 673](#).

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. Both link-local and global unicast address types can either be dynamically assigned (using the Configure Interface page) or manually configured (using the Add IPv6 Address page).
- ◆ An IPv6 global unicast or link-local address can be manually configured (using the Add IPv6 Address page), or a link-local address can be dynamically generated (using the Configure Interface page).

Configuring the IPv6 Default Gateway	Use the IP > IPv6 Configuration (Configure Global) page to configure an IPv6 default gateway for the switch.

Parameters

These parameters are displayed:


- ◆ **Default Gateway** – Sets the IPv6 address of the default next hop router to use when no routing information is known about an IPv6 address.
 - If no static routes are defined, you must define a gateway if the target device is located in a different subnet.
 - 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 456: Configuring the IPv6 Default Gateway



The screenshot shows a web interface for configuring IPv6 settings. At the top, there is a breadcrumb trail 'IP > IPv6'. Below it, an 'Actions:' dropdown menu is set to 'Configure Global'. Underneath, the 'Default Gateway' label is followed by a text input field containing the address '2001:080:2222::254'. At the bottom right of the form, there are two buttons: 'Apply' and 'Revert'.

Configuring IPv6 Interface Settings

Use the IP > IPv6 Configuration (Configure Interface – VLAN or RA Guard) page to configure general IPv6 settings for the selected VLAN, including auto-configuration of a global unicast interface address, and 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 switch's address auto-configuration function will automatically create a link-local address, as well as an IPv6 global address if router advertisements are detected on the local interface.
- ◆ The option to explicitly enable IPv6 creates a link-local address, but will not generate a global IPv6 address if auto-configuration is not enabled. In this case, you can manually configure a global unicast address (see [“Configuring an IPv6 Address” on page 684](#)).
- ◆ The “managed-address configuration” flag tells hosts that they should use stateless address autoconfiguration to get IPv6 address (based on the IPv6 prefixes found in router advertisements) and stateful autoconfiguration to get other non-address parameters (such as DNS server addresses) from DHCPv6 servers. This flag is only a suggestion to attached hosts. They may still use stateful and/or stateless address autoconfiguration. If hosts must be forced to use DHCPv6 for security reasons, ensure that no route prefixes are sent in router advertisements.
- ◆ The “other-stateful-configuration” flag tells hosts that they should use stateful autoconfiguration to obtain information other than addresses from a DHCPv6

server. Some hosts interpret the “other stateful configuration” flag to indicate that they should use stateless address autoconfiguration to get IPv6 address (based on the IPv6 prefixes found in router advertisements) and stateful autoconfiguration to get other non-address parameters from DHCPv6 servers. In this case, the absence of both the “managed address configuration” flag and the “other stateful configuration” flag is interpreted to mean that they should use only stateless autoconfiguration to obtain addresses.

- ◆ 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.
- ◆ **Enable ND managed-config-flag** – Configures IPv6 router advertisements to indicate to attached hosts that they can use stateful autoconfiguration to obtain addresses. (Default: Disabled)
- ◆ **Enable ND other-config-flag** – Configures IPv6 router advertisements to indicate to attached hosts that they can obtain stateful autoconfiguration information other than addresses. (Default: Disabled)
- ◆ **Enable ND ra suppress** – Suppresses periodic unsolicited router advertisements on an interface. It does not suppress advertisements sent in response to a router solicitation. (Default: Disabled)
- ◆ **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: 0-255, 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 165](#)). 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 457: Configuring General Settings for an IPv6 Interface

IP > IPv6 Configuration

Step: 2. Configure Interface

Mode: ☒ VLAN ☐ RA Guard

VLAN: 1

Enable IPv6 Explicitly: ☒ Enabled

Enable ND managed-config-flag: ☐ Enabled

Enable ND other-config-flag: ☐ Enabled

Enable ND ra suppress: ☐ Enabled

MTU (1280-65535): 1500 bytes

ND DAD Attempts (0-255): 1

ND NS Interval (1000-3600000): 1000 ms

ND Reachable-Time (0-3600000): 30000 ms

Apply Reset

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 458: Configuring RA Guard for an IPv6 Interface

IP > IPv6 Configuration

Action: Configure Interface

Mode: ☐ VLAN ☒ RA Guard

Interface: ☒ Port ☐ Trunk

Port List: Total: 10

Port	RA Guard
1	☐ Enabled
2	☐ Enabled
3	☐ Enabled
4	☐ Enabled
5	☐ Enabled

Configuring IPv6 Neighbor Addresses

Use the IP > IPv6 Configuration > Configure Neighbor page to configure a static entry in the IPv6 neighbor discovery cache.

Usage Guidelines

- ◆ Address Resolution Protocol (ARP) has been replaced in IPv6 with the Neighbor Discovery Protocol (NDP). Configuring IPv6 neighbor addresses is similar to configuring static MAC addresses that are implemented using ARP.
- ◆ Static entries can only be configured on an IPv6-enabled interface.
- ◆ The switch does not determine whether a static entry is reachable before placing it in the IPv6 neighbor discovery cache.
- ◆ If the specified entry was dynamically learned through the IPv6 neighbor discovery process, and already exists in the neighbor discovery cache, it is converted to a static entry. Static entries in the IPv6 neighbor discovery cache are not modified if subsequently detected by the neighbor discovery process.
- ◆ Disabling IPv6 on an interface deletes all dynamically learned entries in the IPv6 neighbor discovery cache for that interface, but does not delete static entries.

Parameters

These parameters are displayed:

- ◆ **IPv6 Address** - The IPv6 address of a neighbor device that can be reached through one of the network interfaces configured on this switch. You can specify either a link-local or global unicast address 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.
- ◆ **VLAN** - VLAN ID (Range: 1-4094)
- ◆ **MAC Address** - The 48-bit MAC layer address for the neighbor device. This address must be formatted as six hexadecimal pairs separated by hyphens.

Web Interface

To configure an IPv6 address:

1. Click IP, IPv6 Configuration.
2. Select Configure Neighbor from the Step list.
3. Select Add from the Action list.
4. Enter the IPv6 address and MAC address, then specify the VLAN ID.
5. Click Apply.

Figure 459: Configuring an IPv6 Neighbor Address

The screenshot shows a web interface titled "IP > IPv6 Configuration". At the top, there's a "Step:" dropdown menu set to "3. Configure Neighbor" and an "Action:" dropdown menu set to "Add". Below these are three input fields: "IPv6 Address" (a long text box), "VLAN" (a dropdown menu showing "1"), and "MAC Address" (a text box with a placeholder "(xx-xx-xx-xx-xx-xx or 000000000000)"). At the bottom right, there are two buttons: "Apply" and "Revert".

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 always be configured with a link-local address. Therefore any configuration process that enables IPv6 functionality, or assigns a global unicast address to the switch, including address auto-configuration or explicitly enabling IPv6 (see ["Configuring IPv6 Interface Settings" on page 678](#)), 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 678](#)).
 - 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)
- ◆ **IPv6 Address Mode** – Defines the IPv6 address mode for the VLAN interface.
 - **User Specified** – Configures a specific IPv6 address based on the Address Type setting.
 - **AUTOCONFIG** – Enables stateless autoconfiguration of IPv6 addresses on an interface. The network portion of the address is based on prefixes received in IPv6 router advertisement messages; the host portion is based on the modified EUI-64 form of the interface identifier (i.e., the switch's MAC address).
 - **DHCP** – IPv6 address for the interface must be obtained from a DHCPv6 server.
- ◆ **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 and select the address mode. For the User Specified mode, also select the type and then enter an IPv6 address and prefix length.
4. Click Apply.

Figure 460: Configuring an IPv6 Address

The screenshot shows the 'IP > IPv6 Configuration' web interface. At the top, there is a breadcrumb trail 'IP > IPv6 Configuration'. Below it, a 'Step' dropdown menu is set to '4. Add IPv6 Address'. The main configuration area contains several fields: 'VLAN' with a dropdown menu showing '1', 'IPv6 Address Mode' with a dropdown menu showing 'User Specified', 'IPv6 Address Type' with a dropdown menu showing 'Global', and 'IPv6 Address' with a text input field. At the bottom right of the form, there are two buttons: 'Apply' and 'Reset'.

Configuring IPv6 ND Prefixes

Use the IP > IPv6 Configuration > Configure IPv6 ND Prefix (Add) page to configure IPv6 prefixes to include in router advertisements.

Usage Guidelines

- ◆ Prefixes configured as addresses on an interface are advertised in router advertisements. If prefixes are configured for advertisement using this page, then only these prefixes are advertised.
- ◆ The preferred lifetime and valid lifetime are counted down in real time. After the preferred lifetime expires, no new connections are made using this prefix. When the valid lifetime expires, this prefix will no longer be advertised.
- ◆ All prefixes are inserted in the routing table as Connected (i.e., on-line), unless specified with the off-link option. If the off-link option is specified, and the prefix is already present in the routing table as a Connected prefix, it will be removed.
- ◆ Do not include the link-local prefix in the list of advertised prefixes.

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. (Range: 1-4094)
- ◆ **IPv6 Prefix** – An IPv6 address network prefix.
- ◆ **Default Value** – Uses default values for remaining parameters or can be set to User Specified.
- ◆ **Valid Lifetime** – The amount of time that the specified IPv6 prefix is advertised as being valid. (Range: 0-4294967295 seconds)
- ◆ **Preferred Lifetime** – The amount of time that the specified IPv6 prefix is advertised as being preferred. The preferred lifetime is counted down in real time. (Range: 0-4294967295 seconds)
- ◆ **No-autoconfig** – Indicates to hosts on the local link that the specified prefix cannot be used for IPv6 autoconfiguration.
- ◆ **Off-link** – Indicates that the specified prefix is assigned to the link. Nodes sending traffic to addresses that contain the prefix consider the destination to be locally reachable on the link.

Web Interface

To configure an IPv6 address:

1. Click IP, IPv6 Configuration.
2. Select Configure IPv6 ND Prefix from the Step list.
3. Select Add from the Action list.
4. Specify the VLAN and enter an IPv6 prefix. Set other parameters as needed.
5. Click Apply.

Figure 461: Configuring IPv6 ND Prefixes

The screenshot shows the 'IP > IPv6 Configuration' page. At the top, it says 'Step: 5. Configure IPv6 nd prefix' and 'Action: Add'. Below this, there are several configuration fields: 'VLAN' with a dropdown set to '1', 'IPv6 Prefix' with an empty text box, 'Default Value' with a dropdown set to 'User Specified', 'Valid lifetime' with a text box containing '2592000', 'Preferred lifetime' with a text box containing '604800', 'No-autoconfig' with an unchecked checkbox, and 'Off-link' with an unchecked checkbox. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 462: 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 41: 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.

Table 41: Show IPv6 Neighbors - display description (Continued)

Field	Description
State	The following states are used for dynamic entries: ◆ 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. The following states are used for static entries: ◆ 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 463: Showing IPv6 Neighbors

The screenshot shows the 'IP > IPv6 Configuration' page. Under the 'Action:' dropdown, 'Show IPv6 Neighbor Cache' is selected. Below this, a table titled 'Current Neighbor Cache Table' displays one entry. The table has five columns: IPv6 Address, Age, Link-layer Address, State, and VLAN. The entry shows the IPv6 address 2001:DB8:2222:7272::74, an age of 20, a link-layer address of SE-ED-8C-8C-CA-10, a state of Reachable, and a VLAN of 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	SE-ED-8C-8C-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 42: 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.

Table 42: Show IPv6 Statistics - display description (Continued)

Field	Description
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.
<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.

Table 42: Show IPv6 Statistics - display description (Continued)

Field	Description
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 ipv6IcmpInErrors. 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.
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.

Table 42: Show IPv6 Statistics - display description (Continued)

Field	Description
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.
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 464: Showing IPv6 Statistics (IPv6)

IP > IPv6

Action:

Type: ☒ IPv6 ☐ ICMPv6 ☐ UDP

IPv6 Statistics

Total Received	55	Received Reassembled Succeeded
Received Header Errors	0	Received Reassembled Failed
Received Too Big Errors	0	Transmitted Forwards Datagrams
Received No Routes	0	Transmitted Requests
Received Address Errors	0	Transmitted Discards
Received Unknown Protocols	0	Transmitted No Routes
Received Truncated Packets	0	Transmitted Generated Fragments
Received Discards	0	Transmitted Fragment Succeeded
Received Delivers	55	Transmitted Fragment Failed
Received Reassembly Request Datagrams	0	

Figure 465: Showing IPv6 Statistics (ICMPv6)

IP > IPv6 Configuration

Action:

Type: ☐ IPv6 ☒ ICMPv6 ☐ UDP

ICMPv6 Statistics

Received Input	0	Transmitted Output	0
Received Errors	0	Transmitted Destination Unreachable Messages	0
Received Destination Unreachable Messages	0	Transmitted Packet Too Big Messages	0
Received Packet Too Big Messages	0	Transmitted Time Exceeded Messages	0
Received Time Exceeded Messages	0	Transmitted Parameter Problem Messages	0
Received Parameter Problem Messages	0	Transmitted Echo Request Messages	0
Received Echo Request Messages	0	Transmitted Echo Reply Messages	0
Received Echo Reply Messages	0	Transmitted Router Solicit Messages	0
Received Router Solicit Messages	0	Transmitted Router Advertisement Messages	0
Received Router Advertisement Messages	0	Transmitted Neighbor Solicit Messages	0
Received Neighbor Solicit Messages	0	Transmitted Neighbor Advertisement Messages	0
Received Neighbor Advertisement Messages	0	Transmitted Redirect Messages	0
Received Redirect Messages	0	Transmitted Group Membership Query Messages	0
Received Group Membership Query Messages	0	Transmitted Group Membership Response Messages	0
Received Group Membership Response Messages	0	Transmitted Group Membership Reduction Messages	0
Received Group Membership Reduction Messages	0	Transmitted Multicast Listener Discovery Version 2 Reports	0
Received Multicast Listener Discovery Version 2 Reports	0		

Figure 466: Showing IPv6 Statistics (UDP)

The screenshot shows the 'IP > IPv6' configuration page. Under the 'Actions' dropdown, 'Show Statistics' is selected. The 'Type' section has three radio buttons: 'IPv6', 'ICMPv6', and 'UDP', with 'UDP' being selected. Below this, the 'UDP Statistics' are displayed in a table:

UDP Statistics	
Input	10
No Port Errors	0
Other Errors	0
Output	1

A 'Clear' button is located at the bottom right of the statistics section.

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 43: 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 467: Showing Reported MTU Values

The screenshot shows the 'IP > IPv6 Configuration' page with 'Show MTU' selected in the 'Actions' dropdown. Below the dropdown, it says 'MTU Table Total: 2'. A table displays the reported MTU values:

MTU	Since	Destination Address
1400	00:04:21	5000::1:3
1200	00:04:50	FE80::203:ADFF:FE06:1#1D

General IP Routing

This chapter provides information on network functions including:

- ◆ [Static Routes](#) – Configures static routes to other network segments.
- ◆ [Routing Table](#) – Displays routing entries learned through statically configured entries.


Overview

This switch supports IP routing and routing path management via static routing definitions. 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 routing 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 165](#)), assign all ports that belong to the same group to these VLANs ([page 167](#)), and then assign an IP interface to each VLAN ([page 673](#) or [page 677](#)). 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.

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 468: 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 hop 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:

- ◆ Updating the routing table
- ◆ Updating the Layer 3 switching database

Routing Protocols The switch supports static 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.

Configuring Static Routes

You can enter static routes in the routing table using the IP > Routing > Static Routes (Add) page. Static routes may be required to force the use of a specific route to a subnet. 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.
- ◆ If more than one static routes have the same lowest cost, the first route stored in the routing table will be used.

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 other routing information. (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 469: Configuring Static Routes

IP > Routing > Static Routes

Action: Add

Destination IP Address:

Net Mask / Prefix Length:

Next Hop:

Distance (1-255): (Optional)

Apply Revert

To display static routes:

1. Click IP, Routing, Static Routes.
2. Select Show from the Action List.

Figure 470: Displaying Static Routes

IP > Routing > Static Routes

Action: Show

Static Table List Total: 2

	Destination IP Address	Net Mask	Next Hop	Distance
☐	10.2.48.0	255.255.255.0	10.2.48.1	1
☐	18.5.0.0	255.255.0.0	10.5.38.1	2

Delete Revert

Displaying the Routing Table

Use the IP > Routing > Routing Table (Show Information) page to display all routes that can be accessed via local network interfaces through static routes. If route information is available through more than one of these methods, the priority for route selection is local and then static. 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 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:


- ◆ **Interface** – 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, Others)

Web Interface

To display the routing table:

1. Click IP, Routing, Routing Table.
2. Select Show Information from the Action List.

Figure 471: Displaying the Routing Table



The screenshot shows a web interface titled "IP > Routing > Routing Table". Below the title is an "Action:" dropdown menu set to "Show Information". Below that is a table titled "Routing Table List Total: 2". The table has six columns: "Interface", "Destination IP Address", "Net Mask / Prefix Length", "Next Hop", "Metric", and "Protocol". There are two rows of data.

Interface	Destination IP Address	Net Mask / Prefix Length	Next Hop	Metric	Protocol
lo	127.0.0.1	255.0.0.0	—	0	Local
lo	::1	:::0	—	0	Local

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.
- ◆ **Multicast DNS** – Configures multicast DNS host name-to-address mapping on the local network without the need for a dedicated DNS server.
- ◆ **DHCP** – Configures client, relay, dynamic provisioning, and DHCP server.
- ◆ **DHCPv6 Relay** – Enables the switch as a DHCPv6 Relay Agent and configures destination addresses or VLAN IDs to which client DHCPv6 messages are forwarded.
- ◆ **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 708](#)).

- ◆ 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 472: Configuring General Settings for DNS

The screenshot shows a web interface for configuring DNS settings. The breadcrumb path is 'IP Service > DNS > General'. Below this, there is an 'Action:' dropdown menu currently set to 'Configure Global'. Underneath, the 'Domain Lookup' option is checked with a green checkmark and the word 'Enabled' next to it. Below that, the 'Default Domain Name' is set to 'my.site.com' in a text input field. 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 705](#)). 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 708).
- ◆ 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: 1-127 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 473: Configuring a List of Domain Names for DNS

IP Service > DNS > General

Action: Add Domain Name

Domain Name: sample.com.ak

Apply Revert

To show the list domain names:

1. Click IP Service, DNS.
2. Select Show Domain Names from the Action list.

Figure 474: 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
☐	lind.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 705](#)).
- ◆ 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 475: 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 476: Showing the List of Name Servers for DNS



Configuring Static DNS Host to Address Entries

Use the IP Service > DNS - Static Host Table (Add) page to manually configure static entries in the DNS table that are used to map domain names to IP addresses.

Command Usage

Static entries may be used for local devices connected directly to the attached network, or for commonly used resources located elsewhere on the network.

Parameters

These parameters are displayed:

- ◆ **Host Name** – Name of a host device that is mapped to one or more IP addresses. (Range: 1-127 characters)
- ◆ **IP Address** – IPv4 or IPv6 address(es) associated with a host name.

Web Interface

To configure static entries in the DNS table:

1. Click IP Service, DNS, Static Host Table.
2. Select Add from the Action list.
3. Enter a host name and the corresponding address.
4. Click Apply.

Figure 477: Configuring Static Entries in the DNS Table

IP Service > DNS > Static Host Table

Action: Add

Host Name: yahoo.com

IP Address: 10.2.78.3

Apply Revert

To show static entries in the DNS table:

1. Click IP Service, DNS, Static Host Table.
2. Select Show from the Action list.

Figure 478: Showing Static Entries in the DNS Table

IP Service > DNS > Static Host Table

Action: Show

IP Address List Total: 3

	Host	IP Address
☐	yahoo.com	10.2.78.3
☐	hinet.net	24.29.31.155
☐	google.com	135.48.211.18

Delete Revert

Displaying the DNS Cache Use the IP Service > DNS - Cache page to display entries in the DNS cache that have been learned via the designated name servers.

Command Usage

Servers or other network devices may support one or more connections via multiple IP addresses. If more than one IP address is associated with a host name via information returned from a name server, a DNS client can try each address in succession, until it establishes a connection with the target device.

Parameters

These parameters are displayed:

- ◆ **No.** – The entry number for each resource record.
- ◆ **Flag** – The flag is always “4” indicating a cache entry and therefore unreliable.
- ◆ **Type** – This field includes CNAME which specifies the host address for the owner, and ALIAS which specifies an alias.
- ◆ **IP** – The IP address associated with this record.

- ◆ **TTL** – The time to live reported by the name server.
- ◆ **Host** – The host name associated with this record.

Web Interface

To display entries in the DNS cache:

1. Click IP Service, DNS, Cache.

Figure 479: Showing Entries in the DNS Cache



The screenshot shows a web interface titled "IP Service > DNS > Cache". Below the title, it says "Cache Information: Total: 3". There is a table with 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	A	CNAME	192.168.110.2	360	www.sina.com.cn
2	A	CNAME	10.2.44.3	360	www.yahoo.akadns.net
3	A	ALIAS	portato:2	288	www.yahoo.com

Multicast Domain Name Service

Use the IP Service > Multicast DNS page to enable multicast DNS host name-to-address mapping on the local network without the need for a dedicated DNS server.

Command Usage

- ◆ Multicast DNS allows a network device to choose a domain name in the local DNS name space and announce it using a special multicast IP address. This allows any user to give their computers a link-local mDNS host name of the form "single-dns-label.local." Any name ending in ".local." is therefore link-local, and names within this domain are meaningful only on the link where they originate.

- ◆ When looking for the given host's IP address, the client sends a single-shot mDNS IP multicast query message to all the hosts sharing its local network. Any DNS query for a name ending with ".local." is sent to the mDNS multicast address 224.0.0.251 (or its IPv6 equivalent FF02::FB).

The corresponding host replies with a multicast message announcing itself. All machines in the subnet can then update their mDNS cache with the host's information sent in the reply message.

- ◆ To maintain an on-going cache of host names requires a process of continuous multicast DNS querying. This is done in several phases:
 - Probing – The DNS responder sends a probe message to the local network in order to verify that each entry its local cache is unique.

- **Announcing** – The responder sends an unsolicited mDNS Response containing all of its newly registered resource records (both shared records, and unique records that have completed the probing step).
- **Updating** – The responder repeats the Announcing step to update neighbor caches when the data for any local mDNS record changes.

Parameters

These parameters are displayed:

- ◆ **Multicast DNS Status** – Enables multicast DNS host name-to-address mapping on the local network. (Default: Enabled)

Web Interface

To configure multicast DNS:

1. Click IP Service, Multicast DNS.
2. Mark the check box to enable or disable mDNS as required
3. Click Apply.

Figure 480: Configuring Multicast DNS



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 44: 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 45: 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 481: Specifying a DHCP Client Identifier



IP Service > DHCP > Client

VLAN: 1

Vendor Class ID: ☒ Default

Apply Reset


Configuring DHCP L3 Relay Service

If the switch is configured to provide L3 DHCP Relay service use the IP Service > DHCP > Relay page to configure L3 DHCP relay service for attached host devices. To configure L2 or L3 DHCP relay service refer to the CLI Reference Guide DHCP Relay settings of the CLI Reference Guide - specifically the command: *ip dhcp l2/l3 relay*.

If L3 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 482: 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.

If any of the specified DHCP server addresses are not located in the same network segment with this switch, specify the default router through which this switch can reach other IP subnetworks using the [IP > Routing > Static Routes \(Add\)](#) page (see ["Configuring Static Routes" on page 702](#)) or the [IP > IPv6 Configuration \(Configure Global\)](#) page (see ["Configuring the IPv6 Default Gateway" on page 677](#)).

- ◆ 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 (just below Client)
2. Enter up to five IP addresses for DHCP servers or relay servers in order of preference for any VLAN.
3. Click Apply.

Figure 483: Configuring L3 DHCP Relay Service

IP Service > DHCP > L3 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	IP Address	Relay Agent
1	192.168.3.33	0.0.0.0

Click this button to restart DHCP Relay service.

Configuring DHCP Relay Service with Option 82

Use the IP Service > DHCP > Relay Option 82 page to configure DHCP relay service for attached host devices, including DHCP Option 82 information. DHCP provides an option for sending information about its DHCP clients to the DHCP server (specifically, the interface on the relay server through which the DHCP client request was received). Also known as DHCP Relay Option 82, it allows compatible DHCP servers to use this information when assigning IP addresses, or to set other services or policies for clients.

Option 82 information contains information which can identify both the relay agent and the interface through which the DHCP request was received:

- ◆ The DHCP Relay Information Option Remote ID (RID) is the access node identifier – a string used to identify the switch to the DHCP server.
- ◆ The DHCP Relay Information Option Fields are the Option 82 circuit identification fields (CID – including VLAN ID, stack unit, and port). These fields identify the requesting device by indicating the interface through which the relay agent received the request.

If DHCP relay is enabled, and this switch sees a DHCP client request, it inserts its own IP address into the request so that the DHCP server will know the subnet where the client is located. Depending on the selected frame format set for the remote-id, this information may specify the MAC address, IP address, or an arbitrary string for the requesting device (that is, the relay agent in this context).

By default, the relay agent 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 VLAN ID, stack unit, and port. This allows DHCP client-server exchange messages to be forwarded between the server and client without having to flood them onto the entire VLAN.

The switch then 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.

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.

If any of the specified DHCP server addresses are not located in the same network segment with this switch, specify the default router through which this switch can reach other IP subnetworks.

- ◆ DHCP Snooping Information Option 82 (see [page 364](#)) and DHCP Relay Information Option 82 cannot both be enabled at the same time.
- ◆ DHCP request packets received by the switch are handled as follows:
 - If a DHCP relay server has been set on the switch, when the switch receives a DHCP request packet *without* Option 82 information from the management VLAN or a non-management VLAN, it will add Option 82 relay information and the relay agent's address to the DHCP request packet, and then unicast it to the DHCP server.
 - If a DHCP relay server has been set on the switch, when the switch receives a DHCP request packet *with* Option 82 information from the management VLAN or a non-management VLAN, it will process it according to the configured relay information option policy:
 - If the policy is "replace," the DHCP request packet's Option 82 content (the RID and CID sub-option) is replaced with information provided by the switch. The relay agent address is inserted into the DHCP request packet, and the switch then unicasts this packet to the DHCP server.
 - If the policy is "keep," the DHCP request packet's Option 82 content will be retained. The relay agent address is inserted into the DHCP request packet, and the switch then unicasts this packet to the DHCP server.
 - If the policy is "drop," the original DHCP request packet is flooded onto the VLAN which received the packet but is not relayed.

- ◆ DHCP reply packets received by the relay agent are handled as follows:

When the relay agent receives a DHCP reply packet with Option 82 information over the management VLAN, it first ensures that the packet is destined for it.

- If the RID in the DHCP reply packet is not identical with that configured on the switch, the Option 82 information is retained, and the packet is flooded onto the VLAN through which it was received.
- If the RID in the DHCP reply packet matches that configured on the switch, it then removes the Option 82 information from the packet, and sends it on as follows:
 - If the DHCP packet's broadcast flag is on, the switch uses the circuit-id information contained in the Option 82 information fields to identify

the VLAN connected to the requesting client and then broadcasts the DHCP reply packet to this VLAN.

- If the DHCP packet's broadcast flag is off, the switch uses the circuit-id information in Option 82 fields to identify the interface connected to the requesting client and unicasts the reply packet to the client.
- ◆ DHCP packets are flooded onto the VLAN which received them if DHCP relay service is enabled on the switch and any of the following situations apply:
 - There is no DHCP relay server set on the switch, when the switch receives a DHCP packet.
 - A DHCP relay server has been set on the switch, when the switch receives a DHCP request packet with a non-zero relay agent address field (that is not the address of this switch).
 - A DHCP relay server has been set on the switch, when the switch receives DHCP reply packet without Option 82 information from the management VLAN.
 - The reply packet contains a valid relay agent address field (that is not the address of this switch), or receives a reply packet with a zero relay agent address through the management VLAN.
 - A DHCP relay server has been set on the switch, and the switch receives a reply packet on a non-management VLAN.

Parameters

These parameters are displayed:

- ◆ **Insertion of Relay Information** – Enable DHCP Option 82 information relay. (Default: Disabled)
- ◆ **DHCP Option Policy** – Specifies how to handle client requests which already contain DHCP Option 82 information:
 - **Drop** - Floods the original request packet onto the VLAN that received it instead of relaying it. (This is the default.)
 - **Keep** - Retains the Option 82 information in the client request, inserts the relay agent's address, and unicasts the packet to the DHCP server.
 - **Replace** - Replaces the Option 82 information circuit-id and remote-id fields in the client's request with information provided by the relay agent itself, inserts the relay agent's address, and unicasts the packet to the DHCP server.
- ◆ **DHCP Sub-option Format** – Specifies whether or not to use the sub-type and sub-length fields in the circuit-ID (CID) and remote-ID (RID) in Option 82 information. (Default: Included)

- ◆ **DHCP Remote ID Sub-option** – Sets the remote ID to the switch’s IP address, MAC address, or a user-specified string. (Default: MAC Address (Hex Encoded))
 - **IP Address (Hex/ASCII encoded)** - Inserts an IP address in the remote ID sub-option for the DHCP relay (that is, the IP address of the switch’s CPU).
 - **MAC Address (Hex/ASCII encoded)** - Inserts a MAC address in the remote ID sub-option for the DHCP relay (that is, the MAC address of the switch’s CPU).
 - **String** - An arbitrary string inserted into the remote identifier field. (Range: 1-32 characters)

Web Interface

To configure DHCP relay service:

1. Click IP Service, DHCP, Relay Option 82.
2. Enable or disable the Option 82 relay information.
3. Set the Option 82 policy to specify how to handle Option 82 information already contained in DHCP client request packets.
4. Specify whether or not include “type” and “length” sub-options.
5. Select how to set the Remote ID sub-option.
6. Click Apply.

Figure 484: Configuring DHCP Relay Information Option 82 Service

The screenshot shows a web interface for configuring DHCP Relay Option 82. The breadcrumb trail at the top is 'IP Service > DHCP > Relay Option 82'. The main content area is titled 'Information' and contains four configuration items: 'Insertion of Relay Information' with a checkbox set to 'Enabled', 'DHCP Option Policy' with a dropdown menu set to 'Drop', 'DHCP Sub-option Format' with a dropdown menu set to 'Extra Subtype Included', and 'DHCP Remote ID Sub-option' with a dropdown menu set to 'MAC Address (Hex Encoded)'. At the bottom of the form are 'Apply' and 'Revert' buttons.

Enabling DHCP Dynamic Provision

Use the IP Service > DHCP > Dynamic Provision to enable dynamic provisioning via DHCP.

Command Usage

DHCPD is the daemon used by Linux to dynamically configure TCP/IP information for client systems. To support DHCP option 66/67, you have to add corresponding statements to the configuration file of DHCPD. Information on how to complete

this process are described in [“Downloading a Configuration File and Other Parameters Provided by a DHCP Server”](#) as described in the *CLI Reference Guide*.

Some alternative commands which can be added to the DHCPD to complete the dynamic provisioning process are also described under the **ip dhcp dynamic-provision** command in the *CLI Reference Guide*.

By default, the parameters for DHCP option 66/67 are not carried by the reply sent from the DHCP server. To ask for a DHCP reply with option 66/67, the client can inform the server that it is interested in option 66/67 by sending a DHCP request that includes a 'parameter request list' option. Besides this, the client can also send a DHCP request that includes a 'vendor class identifier' option to the server so that the DHCP server can identify the device, and determine what information should be given to requesting device.

Parameters

These parameters are displayed:

- ◆ **Dynamic Provision via DHCP Status** – Enables dynamic provisioning via DHCP. (Default: Disabled)

Web Interface

To enable dynamic provisioning via DHCP:

1. Click IP Service, DHCP, Dynamic Provision.
2. Mark the Enable box if dynamic provisioning is configured on the DHCP daemon, and required for bootup.
3. Click Apply.

Figure 485: Enabling Dynamic Provisioning via DHCP




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 486: 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 487: Enabling the DHCP Server



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 488: Configuring Excluded Addresses on the DHCP Server

IP Service > DHCP > Server

Step: 2. Configure Excluded Address Actions: Add

Start IP Address: 10.1.0.250

End IP Address: 10.1.0.254 (optional)

Apply Revert

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 489: Showing Excluded Addresses on the DHCP Server

IP Service > DHCP > Server

Step: 2. Configure Excluded Address Actions: Show

DHCP-Excluded Address List Task: 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. 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. Input a Pool Name and click Apply.
5. Select Configure Detail from the Action list.
6. Set the pool Type to Network or Host.
7. 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.
8. Click Apply.

Figure 490: Configuring DHCP Server Address Pools (Network)

The screenshot shows the configuration page for a DHCP server address pool. The breadcrumb is "IP Service > DHCP > Server". The "Step" is "3. Configure Pool" and the "Action" is "Configure Detail".

Pool Name: ipn
Type: Network
IP: 10.1.0.0
Subnet Mask: 255.255.255.0

<<Options>>

Default Router	10.1.0.253	Default Router 2	0.0.0.0
DNS Server	10.2.3.4	DNS Server 2	0.0.0.0
Netbios Server	10.1.0.33	Netbios Server 2	0.0.0.0
Netbios Type	Hybrid	Domain Name	example.com
Bootfile	wire.boot	Next Server	10.1.0.21
Lease Time	Infinite		

Buttons: Apply, Revert

Figure 491: Configuring DHCP Server Address Pools (Host)

The screenshot shows the configuration page for a DHCP server address pool. The breadcrumb is "IP Service > DHCP > Server". The "Step" is "3. Configure Pool" and the "Action" is "Configure Detail".

Pool Name: mgr
Type: Host
IP: 10.1.0.19
Subnet Mask: 255.255.255.0

Client Identifier: Text: boar
Hardware Address: Ethernet: 00-e0-29-34-26 (XX-XX-XX-XX-XX-XX OF XXXXXXXXXXXX)

<<Options>>

Default Router	10.1.0.253	Default Router 2	0.0.0.0
DNS Server	10.2.3.4	DNS Server 2	0.0.0.0
Netbios Server	0.0.0.0	Netbios Server 2	0.0.0.0
Netbios Type	Hybrid	Domain Name	example.com
Bootfile	wire.boot	Next Server	10.1.0.21
Lease Time	Infinite		

Buttons: Apply, Revert

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 492: Showing Configured DHCP Server Address Pools



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 493: Shows Addresses Assigned by the DHCP Server



Configuring DHCPv6 Relay

Use IP Service > DHCPv6 > Relay (Add) page to enable the switch as a DHCPv6 Relay Agent and configure destination addresses or VLAN IDs to which client DHCPv6 messages are forwarded. When the relay agent is enabled on a specified VLAN, the switch listens to UDP port 547 for DHCPv6 messages on that VLAN. The agent will forward the messages to the configured unicast addresses or to a multicast address on configured VLANs or all VLANs.

Command Usage

- ◆ When the unicast addresses of the DHCPv6 servers or other relay agents are known, configure the addresses for each VLAN that will use the IPv6 DHCP relay service.
- ◆ Up to five DHCPv6 server or relay agent unicast addresses can be configured.
- ◆ Use the “Multicast mode”, if you do not know the Unicast destination address of the DHCPv6 servers or other relay agents.
- ◆ All associated VLANs must have their L3 interface enabled using the VLAN > Static (Action:Modify) page.
- ◆ The relay agent is enabled when at least one configured VLAN has an entry is listed in the IP Service > DHCPv6 > Relay (Action: Show) page.
- ◆ The relay agent is disabled if there are no entries for all configured VLANs in the IP Service > DHCPv6 > Relay (Action: Show) page.

Parameters

These parameters are displayed:

- ◆ **VLAN** – ID of the configured VLAN.
- ◆ **Mode:**
 - **Unicast** - To forward the client IPv6 DHCP messages to a specific destination address(es), select the Unicast mode.
 - **Multicast** - To forward the client IPv6 DHCP messages to all available servers on configured VLANs, select the Multicast mode. In this mode the DHCPv6 agent forwards DHCPv6 messages to the IPv6 DHCP multicast address known as “All_DHCP_Servers” (FF05::1:3) which both DHCPv6 servers and relay agents listen to.
- ◆ **Destination IP** - When Unicast mode is selected, enter an IPv6 address of a DHCPv6 server or relay agent (Maximum: 5 for each configured VLAN).

- ◆ **Destination VLAN** – When Multicast mode is selected, select a configured VLAN or “All” (all VLANs). The DHCPv6 messages will be forwarded on the configured VLANs or all VLANs using the “All_DHCP_Servers” address.

Web Interface

To enable the switch’s DHCPv6 Relay Agent for Unicast mode:

1. Click IP Service, DHCPv6, Relay (Action: Add).
2. Select a VLAN from the VLAN drop-down list.
3. Select Unicast from the Mode drop-down list.
4. Input the server or relay agent’s IPv6 address in the destination IP field.
5. Click Apply

Figure 494: Enabling DHCPv6 Relay Agent for Unicast mode.

The screenshot shows the configuration page for the DHCPv6 Relay Agent in Unicast mode. The breadcrumb navigation at the top reads "IP Service > DHCPv6 > Relay". Below this, there is a section for configuration with the following fields:

- Action:** A dropdown menu set to "Add".
- VLAN:** A dropdown menu set to "1".
- Mode:** A dropdown menu set to "Unicast".
- Destination IP:** A text field containing the IPv6 address "2001:0db8:05e3:0000:0000:002e:0370:7334".

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

To enable the switch’s DHCPv6 Relay Agent for Multicast mode:

1. Click IP Service, DHCPv6, Relay (Action: Add).
2. Select a VLAN from the VLAN drop-down list.
3. Select Multicast from the Mode drop-down list.
4. Select a destination VLAN or All (VLANs) from the Destination VLAN drop-down list.
5. Click Apply

By default, the parameters for DHCPv6 option 59/60 are not carried by the reply sent from the DHCPv6 server. To ask for a DHCPv6 reply with option 59/60, the client can inform the server that it is interested in option 59/60 by sending a DHCPv6 information request.

Parameters

These parameters are displayed:

- ◆ **Dynamic Provision via DHCPv6 Status** – Enables dynamic provisioning via DHCP. (Default: Disabled)

Web Interface

To enable dynamic provisioning via DHCPv6:

1. Click IP Service, DHCPv6, Dynamic Provision.
2. Mark the Enable box if dynamic provisioning is configured on the DHCP daemon, and required for bootstrap.
3. Click Apply.

Figure 497: Enabling Dynamic Provisioning via DHCPv6



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 498: Configuring Global Settings for PPPoE Intermediate Agent

IP Service > PPPoE Intermediate Agent

Step: 1. Configure Global

PPPoE IA Global Status: ☐ Enabled

Access Mode Identifier:

Operational Access Mode Identifier: 18.254.254.1

Generic Error Message:

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.

◆ **Operation 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)

◆ **Operation Remote ID** – The configured circuit identifier.

◆ **Remote ID Delimiter** – Enables a user-specified delimiter for the remote ID.

◆ **Delimiter ASCII** – A character used to separate components in the remote circuit ID value. (Range: 0-255; Default: 35, ASCII character "#")

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, set the circuit ID, remote ID, or remote ID delimiter.
5. Click Apply.

Figure 499: Configuring Interface Settings for PPPoE Intermediate Agent

IP Service > PPPoE Intermediate Agent

Step: 2. Configure Interface

Interface: ☒ Port ☐ Trunk

PPPoE Intermediate Agent Port List Total: 20

Port	PPPoE IA Status	Trust Status	Vendor Tag Strip	Circuit ID	Operation Circuit ID	Name ID	Operation Name ID	Remote ID	Delimiter ASCII (0-255)
1	☐ Enabled	☐ Enabled	☐ Enabled		1/1-vtd		10-3C-E3-8C-44-83	☐ Enabled	35
2	☐ Enabled	☐ Enabled	☐ Enabled		1/2-vtd		10-3C-E3-8C-44-84	☐ Enabled	35
3	☐ Enabled	☐ Enabled	☐ Enabled		1/3-vtd		10-3C-E3-8C-44-85	☐ Enabled	35
4	☐ Enabled	☐ Enabled	☐ Enabled		1/4-vtd		10-3C-E3-8C-44-86	☐ Enabled	35
5	☐ Enabled	☐ Enabled	☐ Enabled		1/5-vtd		10-3C-E3-8C-44-87	☐ Enabled	35

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 500: Showing PPPoE Intermediate Agent Statistics

PPPoE Intermediate Agent Statistics Total: 23									
Port	Received						Dropped		
	AE	PAD1	PAD2	PAD3	PAD4	PAD5	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

Section III

Appendices

This section provides additional information and includes these items:

- ◆ [“Software Specifications” on page 739](#)
- ◆ [“Troubleshooting” on page 745](#)
- ◆ [“License Information” on page 747](#)



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
10GBASE-SR/LR/ER: 10 Gbps at full duplex (SFP+)
25GBASE-SR/LR: 25 Gbps at full duplex (SFP28)

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 3 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 eight levels of priority
Strict, Weighted Round Robin (WRR), or a combination of strict and weighted queuing
Layer 3/4 priority mapping: IP DSCP

Quality of Service DiffServ¹⁸ supports class maps, policy maps, and service policies

Multicast Filtering IGMP Snooping (Layer 2 IPv4)
MLD Snooping (Layer 2 IPv6)
Multicast VLAN Registration (IPv4/IPv6)

IP Routing ARP, CIDR (Classless Inter-Domain Routing)

Additional Features BOOTP Client
CFM (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

¹⁸. Only supported for IPv4.

Software Loading HTTP, FTP, SFTP, SCP, TFTP in-band, or XModem out-of-band

SNMP Management access via MIB database
Trap management to specified hosts

RMON Groups 1, 2, 3, 9 (Statistics, History, Alarm, Event)

Standards

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
Link Aggregation Control Protocol (LACP)
Full-duplex flow control (ISO/IEC 8802-3)
IEEE 802.3ac VLAN tagging
ARP (RFC 826)
DHCP Client (RFC 2131)
DHCP Relay (RFC 951, 2132, 3046)
HTTPS
ICMP (RFC 792)
IGMP (RFC 1112)
IGMPv2 (RFC 2236)
IGMP Proxy (RFC 4541)
IPv4 IGMP (RFC 3228)
MLD Snooping (RFC 4541)
NTP (RFC 1305)
RADIUS+ (RFC 2618)
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)

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)
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)

Problems Accessing the Management Interface

Table 46: Troubleshooting Chart

Symptom	Action
Cannot connect using Telnet, web browser, or SNMP software	◆ 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	◆ 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 an RSA 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	◆ 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	◆ Contact your local distributor.

Using System Logs

If a fault does occur, refer to the Installation Guide to ensure that the problem you encountered is actually caused by the switch. If the problem appears to be caused by the switch, follow these steps:

1. Enable logging.
2. Set the error messages reported to include all categories.
3. Enable SNMP.
4. Enable SNMP traps.
5. Designate the SNMP host that is to receive the error messages.
6. Repeat the sequence of commands or other actions that lead up to the error.
7. Make a list of the commands or circumstances that led to the fault. Also make a list of any error messages displayed.
8. Set up your terminal emulation software so that it can capture all console output to a file. Then enter the “show tech-support” command to record all system settings in this file.
9. Contact your distributor’s service engineer, and send a detailed description of the problem, along with the file used to record your system settings.

For example:

```
Console(config)#logging on
Console(config)#logging history flash 7
Console(config)#snmp-server host 192.168.1.23
:
```



License Information

This product includes copyrighted third-party software subject to the terms of the GNU General Public License (GPL), GNU Lesser General Public License (LGPL), or other related free software licenses. The GPL code used in this product is distributed WITHOUT ANY WARRANTY and is subject to the copyrights of one or more authors. For details, refer to the section "The GNU General Public License" below, or refer to the applicable license as included in the source-code archive.

The GNU General Public License

GNU GENERAL PUBLIC LICENSE
Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.
59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

1. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

2. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

3. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
 - a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
 - b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
 - c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

4. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
 - a). Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

- 5. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
- 6. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
- 7. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
- 8. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

9. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
10. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.
11. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

1. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
2. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

