



ECS4120-28F/T/P
ECS4120-52T
28/52-Port Layer 2+
Gigabit Ethernet Switch

Software Release 1.0.2.14

Web Management Guide

Web Management Guide

ECS4120-28T Gigabit Ethernet Switch

L2+ Gigabit Ethernet Switch
with 24 10/100/1000BASE-T (RJ-45) Ports,
and 4 Gigabit SFP Ports

ECS4120-28F Gigabit Ethernet Switch

L2+ Gigabit Ethernet Switch
with 20 100/1000 SFP Ports,
4 10/100/1000 BASE-T (RJ-45) / 100/1000 SFP Combo Ports,
and 4 10 Gigabit SFP+ Ports

ECS4120-28P Gigabit Ethernet Switch

L2+ Gigabit Ethernet Switch
with 24 10/100/1000BASE-T (RJ-45) 802.3 af/at PoE Ports,
and 4 10-Gigabit SFP+ Ports
(PoE Power Budget: 370W)

ECS4120-52T Gigabit Ethernet Switch

L2+ Gigabit Ethernet Switch
with 48 10/100/1000BASE-T (RJ-45) Ports,
and 4 Gigabit SFP Ports

How to Use This Guide

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

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

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

The guide includes these sections:

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

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

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

CLI Reference Guide



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

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

Installation Guide

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

Quick Start Guide

Safety and Regulatory Information

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



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



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

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

November 2016 Revision

This is the first version of this guide. This guide is valid for software release v1.0.2.14.

Contents

How to Use This Guide	3
Contents	5
Figures	17
Tables	31

Section I	Getting Started	33
	1 Introduction	35
	Key Features	35
	Description of Software Features	36
	Address Resolution Protocol	41
	Operation, Administration, and Maintenance	41
	Multicast Filtering	41
	Link Layer Discovery Protocol	41
	System Defaults	42

Section II	Web Configuration	45
	2 Using the Web Interface	47
	Connecting to the Web Interface	47
	Navigating the Web Browser Interface	48
	Dashboard	48
	Home Page	50
	Configuration Options	50
	Panel Display	51
	Main Menu	52

3 Basic Management Tasks	71
Displaying System Information	72
Displaying Hardware/Software Versions	73
Configuring Support for Jumbo Frames	74
Displaying Bridge Extension Capabilities	75
Managing System Files	77
Copying Files via FTP/ TFTP or HTTP	77
Saving the Running Configuration to a Local File	79
Setting the Start-up File	80
Showing System Files	81
Automatic Operation Code Upgrade	81
Setting the System Clock	85
Setting the Time Manually	86
Setting the SNTP Polling Interval	87
Configuring NTP	87
Configuring Time Servers	88
Setting the Time Zone	92
Configuring Summer Time	93
Configuring the Console Port	95
Configuring Telnet Settings	97
Displaying CPU Utilization	98
Configuring CPU Guard	99
Displaying Memory Utilization	100
Resetting the System	101
4 Interface Configuration	105
Port Configuration	106
Configuring by Port List	106
Configuring by Port Range	109
Displaying Connection Status	110
Showing Port or Trunk Statistics	111
Displaying Statistical History	116
Transceiver Data and Thresholds	120
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
VLAN Trunking	156
5 VLAN Configuration	159
IEEE 802.1Q VLANs	159
Configuring VLAN Groups	163
Adding Static Members to VLANs	165
Configuring Dynamic VLAN Registration	170
IEEE 802.1Q Tunneling	173
Enabling QinQ Tunneling on the Switch	177
Creating CVLAN to SPVLAN Mapping Entries	178
Adding an Interface to QinQ Tunnel	180
L2PT Tunneling	181
Configuring the L2PT Tunnel Address	183
Enabling L2PT for Selected Interfaces	184
Protocol VLANs	185
Configuring Protocol VLAN Groups	186
Mapping Protocol Groups to Interfaces	187
Configuring IP Subnet VLANs	189

Configuring MAC-based VLANs	191
Configuring VLAN Translation	193
6 Address Table Settings	197
Dynamic Address Cache	197
Displaying the Dynamic Address Table	197
Clearing the Dynamic Address Table	198
Changing the Aging Time	199
Configuring MAC Address Learning	200
Setting Static Addresses	201
Issuing MAC Address Traps	203
7 Spanning Tree Algorithm	205
Overview	205
Configuring Loopback Detection	208
Configuring Global Settings for STA	209
Displaying Global Settings for STA	215
Configuring Interface Settings for STA	216
Displaying Interface Settings for STA	220
Configuring Multiple Spanning Trees	223
Configuring Interface Settings for MSTP	227
8 Congestion Control	231
Rate Limiting	231
Storm Control	232
Automatic Traffic Control	234
Setting the ATC Timers	236
Configuring ATC Thresholds and Responses	237
9 Class of Service	241
Layer 2 Queue Settings	241
Setting the Default Priority for Interfaces	241
Selecting the Queue Mode	242
Mapping CoS Values to Egress Queues	245
Layer 3/4 Priority Settings	248
Setting Priority Processing to IP Precedence/DSCP or CoS	248
Mapping Ingress DSCP Values to Internal DSCP Values	249

Mapping CoS Priorities to Internal DSCP Values	252
Mapping Internal DSCP Values to Egress CoS Values	254
Mapping IP Precedence Values to Internal DSCP Values	256
Mapping IP Port Priority to Internal DSCP Values	258
10 Quality of Service	261
Overview	261
Configuring a Class Map	262
Creating QoS Policies	265
Attaching a Policy Map to a Port	275
11 VoIP Traffic Configuration	277
Overview	277
Configuring VoIP Traffic	278
Configuring Telephony OUI	279
Configuring VoIP Traffic Ports	280
12 Security Measures	283
AAA Authentication, Authorization and Accounting	284
Configuring Local/Remote Logon Authentication	285
Configuring Remote Logon Authentication Servers	286
Configuring AAA Accounting	291
Configuring AAA Authorization	296
Configuring User Accounts	299
Web Authentication	301
Configuring Global Settings for Web Authentication	301
Configuring Interface Settings for Web Authentication	302
Network Access (MAC Address Authentication)	304
Configuring Global Settings for Network Access	306
Configuring Network Access for Ports	307
Configuring Port Link Detection	309
Configuring a MAC Address Filter	310
Displaying Secure MAC Address Information	312
Configuring HTTPS	313
Configuring Global Settings for HTTPS	313
Replacing the Default Secure-site Certificate	315

Configuring the Secure Shell	316
Configuring the SSH Server	319
Generating the Host Key Pair	320
Importing User Public Keys	322
Access Control Lists	324
Showing TCAM Utilization	325
Setting the ACL Name and Type	327
Configuring a Standard IPv4 ACL	329
Configuring an Extended IPv4 ACL	330
Configuring a Standard IPv6 ACL	333
Configuring an Extended IPv6 ACL	334
Configuring a MAC ACL	336
Configuring an ARP ACL	338
Binding a Port to an Access Control List	340
Showing ACL Hardware Counters	341
Filtering IP Addresses for Management Access	342
Configuring Port Security	344
Configuring 802.1X Port Authentication	347
Configuring 802.1X Global Settings	348
Configuring Port Authenticator Settings for 802.1X	349
Displaying 802.1X Statistics	353
DoS Protection	355
DHCPv4 Snooping	356
DHCP Snooping Global Configuration	359
DHCP Snooping VLAN Configuration	360
Configuring Interfaces for DHCP Snooping	361
Displaying DHCP Snooping Binding Information	363
DHCPv6 Snooping	364
DHCPv6 Snooping Global Configuration	366
DHCPv6 Snooping VLAN Configuration	368
Configuring Interfaces for DHCPv6 Snooping	369
Displaying DHCPv6 Snooping Binding Information	371
Displaying DHCPv6 Snooping Statistics	372
IPv4 Source Guard	373
Configuring Ports for IPv4 Source Guard	373

Configuring Static Bindings for IPv4 Source Guard	375
Displaying Information for Dynamic IPv4 Source Guard Bindings	378
IPv6 Source Guard	379
Configuring Ports for IPv6 Source Guard	379
Configuring Static Bindings for IPv6 Source Guard	381
Displaying Information for Dynamic IPv6 Source Guard Bindings	384
ARP Inspection	385
Configuring Global Settings for ARP Inspection	386
Configuring VLAN Settings for ARP Inspection	388
Configuring Interface Settings for ARP Inspection	389
Displaying ARP Inspection Statistics	391
Displaying the ARP Inspection Log	392
Application Filter	393
13 Basic Administration Protocols	395
Configuring Event Logging	396
System Log Configuration	396
Remote Log Configuration	398
Sending Simple Mail Transfer Protocol Alerts	399
Link Layer Discovery Protocol	401
Setting LLDP Timing Attributes	401
Configuring LLDP Interface Attributes	403
Configuring LLDP Interface Civic-Address	407
Displaying LLDP Local Device Information	409
Displaying LLDP Remote Device Information	413
Displaying Device Statistics	421
Power over Ethernet	423
Setting the Switch's Overall PoE Power Budget	423
Setting The Port PoE Power Budget	424
Simple Network Management Protocol	426
Configuring Global Settings for SNMP	428
Setting Community Access Strings	429
Setting the Local Engine ID	431
Specifying a Remote Engine ID	432
Setting SNMPv3 Views	433

Configuring SNMPv3 Groups	436
Configuring Local SNMPv3 Users	441
Configuring Remote SNMPv3 Users	443
Specifying Trap Managers	446
Creating SNMP Notification Logs	450
Showing SNMP Statistics	452
Remote Monitoring	454
Configuring RMON Alarms	454
Configuring RMON Events	457
Configuring RMON History Samples	459
Configuring RMON Statistical Samples	462
Switch Clustering	464
Configuring General Settings for Clusters	465
Cluster Member Configuration	466
Managing Cluster Members	468
Setting a Time Range	469
Ethernet Ring Protection Switching	471
ERPS Global Configuration	476
ERPS Ring Configuration	476
ERPS Forced and Manual Mode Operations	492
OAM Configuration	496
Enabling OAM on Local Ports	496
Displaying Statistics for OAM Messages	499
Displaying the OAM Event Log	500
Displaying the Status of Remote Interfaces	501
Configuring a Remote Loopback Test	502
Displaying Results of Remote Loopback Testing	504
Connectivity Fault Management	505
Configuring Global Settings for CFM	509
Configuring Interfaces for CFM	512
Configuring CFM Maintenance Domains	512
Configuring CFM Maintenance Associations	517
Configuring Maintenance End Points	521
Configuring Remote Maintenance End Points	523
Transmitting Link Trace Messages	525

Transmitting Loop Back Messages	527
Transmitting Delay-Measure Requests	529
Displaying Local MEPs	531
Displaying Details for Local MEPs	532
Displaying Local MIPs	534
Displaying Remote MEPs	535
Displaying Details for Remote MEPs	536
Displaying the Link Trace Cache	538
Displaying Fault Notification Settings	539
Displaying Continuity Check Errors	540
OAM Configuration	541
Enabling OAM on Local Ports	541
Displaying Statistics for OAM Messages	544
Displaying the OAM Event Log	545
Displaying the Status of Remote Interfaces	546
Configuring a Remote Loopback Test	547
Displaying Results of Remote Loopback Testing	549
UDLD Configuration	550
Configuring UDLD Protocol Intervals	551
Configuring UDLD Interface Settings	552
Displaying UDLD Neighbor Information	554
LBD Configuration	555
Configuring Global Settings for LBD	556
Configuring Interface Settings for LBD	558
14 Multicast Filtering	559
Overview	559
Layer 2 IGMP (Snooping and Query for IPv4)	560
Configuring IGMP Snooping and Query Parameters	562
Specifying Static Interfaces for a Multicast Router	566
Assigning Interfaces to Multicast Services	568
Setting IGMP Snooping Status per Interface	570
Filtering IGMP Query Packets and Multicast Data	576
Displaying Multicast Groups Discovered by IGMP Snooping	577
Displaying IGMP Snooping Statistics	578

Filtering and Throttling IGMP Groups	582
Enabling IGMP Filtering and Throttling	583
Configuring IGMP Filter Profiles	583
Configuring IGMP Filtering and Throttling for Interfaces	586
MLD Snooping (Snooping and Query for IPv6)	587
Configuring MLD Snooping and Query Parameters	588
Setting Immediate Leave Status for MLD Snooping per Interface	589
Specifying Static Interfaces for an IPv6 Multicast Router	590
Assigning Interfaces to IPv6 Multicast Services	592
Showing MLD Snooping Groups and Source List	594
Multicast VLAN Registration for IPv4	596
Configuring MVR Global Settings	597
Configuring MVR Domain Settings	599
Configuring MVR Group Address Profiles	600
Configuring MVR Interface Status	603
Assigning Static MVR Multicast Groups to Interfaces	605
Displaying MVR Receiver Groups	607
Displaying MVR Statistics	608
Multicast VLAN Registration for IPv6	612
Configuring MVR6 Global Settings	613
Configuring MVR6 Domain Settings	615
Configuring MVR6 Group Address Profiles	616
Configuring MVR6 Interface Status	619
Assigning Static MVR6 Multicast Groups to Interfaces	621
Displaying MVR6 Receiver Groups	623
Displaying MVR6 Statistics	624
15 Basic IP Functions	629
Using the Ping Function	629
Using the Trace Route Function	631
Address Resolution Protocol	632
Basic ARP Configuration	633
Configuring Static ARP Addresses	634
Displaying Dynamic or Local ARP Entries	636
Displaying ARP Statistics	636

16 IP Configuration	639
Setting the Switch's IP Address (IP Version 4)	639
Configuring IPv4 Interface Settings	639
Setting the Switch's IP Address (IP Version 6)	643
Configuring the IPv6 Default Gateway	643
Configuring IPv6 Interface Settings	644
Configuring an IPv6 Address	649
Showing IPv6 Addresses	652
Showing the IPv6 Neighbor Cache	653
Showing IPv6 Statistics	654
Showing the MTU for Responding Destinations	660
17 General IP Routing	661
Overview	661
Initial Configuration	661
IP Routing and Switching	662
Routing Path Management	663
Routing Protocols	663
Configuring IP Routing Interfaces	664
Configuring Local and Remote Interfaces	664
Configuring Static Routes	665
Displaying the Routing Table	666
18 IP Services	669
Domain Name Service	669
Configuring General DNS Service Parameters	669
Configuring a List of Domain Names	670
Configuring a List of Name Servers	672
Configuring Static DNS Host to Address Entries	673
Displaying the DNS Cache	674
Dynamic Host Configuration Protocol	675
Specifying a DHCP Client Identifier	675
Configuring DHCP Relay Service	677
Enabling DHCP Dynamic Provision	678
Configuring the PPPoE Intermediate Agent	679
Configuring PPPoE IA Global Settings	679

Configuring PPPoE IA Interface Settings	681
Showing PPPoE IA Statistics	682

Section III	Appendices	685
	A Software Specifications	687
	Software Features	687
	Management Features	689
	Standards	689
	Management Information Bases	690
	B Troubleshooting	693
	Problems Accessing the Management Interface	693
	Using System Logs	694
	C License Information	695
	The GNU General Public License	695
	Glossary	699
	Index	707

Figures

Figure 1: Dashboard	48
Figure 2: Home Page	50
Figure 3: Front Panel Indicators	51
Figure 4: System Information	72
Figure 5: General Switch Information	74
Figure 6: Configuring Support for Jumbo Frames	75
Figure 7: Displaying Bridge Extension Configuration	76
Figure 8: Copy Firmware	78
Figure 9: Saving the Running Configuration	80
Figure 10: Setting Start-Up Files	80
Figure 11: Displaying System Files	81
Figure 12: Configuring Automatic Code Upgrade	85
Figure 13: Manually Setting the System Clock	86
Figure 14: Setting the Polling Interval for SNTP	87
Figure 15: Configuring NTP	88
Figure 16: Specifying SNTP Time Servers	89
Figure 17: Adding an NTP Time Server	90
Figure 18: Showing the NTP Time Server List	90
Figure 19: Adding an NTP Authentication Key	91
Figure 20: Showing the NTP Authentication Key List	92
Figure 21: Setting the Time Zone	93
Figure 22: Configuring Summer Time	95
Figure 23: Console Port Settings	96
Figure 24: Telnet Connection Settings	98
Figure 25: Displaying CPU Utilization	99
Figure 26: Configuring CPU Guard	100
Figure 27: Displaying Memory Utilization	101
Figure 28: Restarting the Switch (Immediately)	103
Figure 29: Restarting the Switch (In)	103

Figure 30: Restarting the Switch (At)	104
Figure 31: Restarting the Switch (Regularly)	104
Figure 32: Configuring Connections by Port List	109
Figure 33: Configuring Connections by Port Range	110
Figure 34: Displaying Port Information	111
Figure 35: Showing Port Statistics (Table)	115
Figure 36: Showing Port Statistics (Chart)	115
Figure 37: Configuring a History Sample	117
Figure 38: Showing Entries for History Sampling	118
Figure 39: Showing Status of Statistical History Sample	118
Figure 40: Showing Current Statistics for a History Sample	119
Figure 41: Showing Ingress Statistics for a History Sample	120
Figure 42: Displaying Transceiver Data	121
Figure 43: Configuring Transceiver Thresholds	123
Figure 44: Performing Cable Tests	125
Figure 45: Configuring Static Trunks	126
Figure 46: Creating Static Trunks	127
Figure 47: Adding Static Trunks Members	128
Figure 48: Configuring Connection Parameters for a Static Trunk	128
Figure 49: Showing Information for Static Trunks	129
Figure 50: Configuring Dynamic Trunks	129
Figure 51: Configuring the LACP Aggregator Admin Key	132
Figure 52: Enabling LACP on a Port	133
Figure 53: Configuring LACP Parameters on a Port	133
Figure 54: Showing Members of a Dynamic Trunk	134
Figure 55: Configuring Connection Settings for a Dynamic Trunk	134
Figure 56: Displaying Connection Parameters for Dynamic Trunks	135
Figure 57: Displaying LACP Port Counters	136
Figure 58: Displaying LACP Port Internal Information	137
Figure 59: Displaying LACP Port Remote Information	139
Figure 60: Configuring Load Balancing	140
Figure 61: Enabling Power Savings	142
Figure 62: Configuring Local Port Mirroring	142
Figure 63: Configuring Local Port Mirroring	143
Figure 64: Displaying Local Port Mirror Sessions	144

Figure 65: Configuring Remote Port Mirroring	144
Figure 66: Configuring Remote Port Mirroring (Source)	147
Figure 67: Configuring Remote Port Mirroring (Intermediate)	148
Figure 68: Configuring Remote Port Mirroring (Destination)	148
Figure 69: Configuring an sFlow Receiver	150
Figure 70: Showing sFlow Receivers	151
Figure 71: Configuring an sFlow Instance	152
Figure 72: Showing sFlow Instances	152
Figure 73: Enabling Traffic Segmentation	154
Figure 74: Configuring Members for Traffic Segmentation	155
Figure 75: Showing Traffic Segmentation Members	156
Figure 76: Configuring VLAN Trunking	156
Figure 77: Configuring VLAN Trunking	157
Figure 78: VLAN Compliant and VLAN Non-compliant Devices	161
Figure 79: Using GVRP	162
Figure 80: Creating Static VLANs	164
Figure 81: Modifying Settings for Static VLANs	165
Figure 82: Showing Static VLANs	165
Figure 83: Configuring Static Members by VLAN Index	168
Figure 84: Configuring Static VLAN Members by Interface	169
Figure 85: Configuring Static VLAN Members by Interface Range	169
Figure 86: Configuring Global Status of GVRP	171
Figure 87: Configuring GVRP for an Interface	172
Figure 88: Showing Dynamic VLANs Registered on the Switch	172
Figure 89: Showing the Members of a Dynamic VLAN	173
Figure 90: QinQ Operational Concept	174
Figure 91: Enabling QinQ Tunneling	178
Figure 92: Configuring CVLAN to SPVLAN Mapping Entries	179
Figure 93: Showing CVLAN to SPVLAN Mapping Entries	179
Figure 94: Adding an Interface to a QinQ Tunnel	181
Figure 95: Configuring the L2PT Tunnel Address	184
Figure 96: Enabling L2PT on Required Interfaces	185
Figure 97: Configuring Protocol VLANs	187
Figure 98: Displaying Protocol VLANs	187
Figure 99: Assigning Interfaces to Protocol VLANs	188

Figure 100: Showing the Interface to Protocol Group Mapping	189
Figure 101: Configuring IP Subnet VLANs	191
Figure 102: Showing IP Subnet VLANs	191
Figure 103: Configuring MAC-Based VLANs	192
Figure 104: Showing MAC-Based VLANs	193
Figure 105: Configuring VLAN Translation	193
Figure 106: Configuring VLAN Translation	194
Figure 107: Showing the Entries for VLAN Translation	195
Figure 108: Displaying the Dynamic MAC Address Table	198
Figure 109: Clearing Entries in the Dynamic MAC Address Table	199
Figure 110: Setting the Address Aging Time	200
Figure 111: Configuring MAC Address Learning	201
Figure 112: Configuring Static MAC Addresses	202
Figure 113: Displaying Static MAC Addresses	203
Figure 114: Issuing MAC Address Traps (Global Configuration)	204
Figure 115: Issuing MAC Address Traps (Interface Configuration)	204
Figure 116: STP Root Ports and Designated Ports	206
Figure 117: MSTP Region, Internal Spanning Tree, Multiple Spanning Tree	207
Figure 118: Spanning Tree – Common Internal, Common, Internal	207
Figure 119: Configuring Port Loopback Detection	209
Figure 120: Configuring Global Settings for STA (STP)	213
Figure 121: Configuring Global Settings for STA (RSTP)	214
Figure 122: Configuring Global Settings for STA (MSTP)	214
Figure 123: Displaying Global Settings for STA	216
Figure 124: Determining the Root Port	217
Figure 125: Configuring Interface Settings for STA	220
Figure 126: STA Port Roles	222
Figure 127: Displaying Interface Settings for STA	223
Figure 128: Creating an MST Instance	224
Figure 129: Displaying MST Instances	225
Figure 130: Modifying the Priority for an MST Instance	225
Figure 131: Displaying Global Settings for an MST Instance	226
Figure 132: Adding a VLAN to an MST Instance	226
Figure 133: Displaying Members of an MST Instance	227
Figure 134: Configuring MSTP Interface Settings	228

Figure 135: Displaying MSTP Interface Settings	229
Figure 136: Configuring Rate Limits	232
Figure 137: Configuring Storm Control	234
Figure 138: Storm Control by Limiting the Traffic Rate	234
Figure 139: Storm Control by Shutting Down a Port	235
Figure 140: Configuring ATC Timers	237
Figure 141: Configuring ATC Interface Attributes	239
Figure 142: Setting the Default Port Priority	242
Figure 143: Setting the Queue Mode (Strict)	244
Figure 144: Setting the Queue Mode (WRR)	244
Figure 145: Setting the Queue Mode (Strict and WRR)	245
Figure 146: Mapping CoS Values to Egress Queues	247
Figure 147: Showing CoS Values to Egress Queue Mapping	247
Figure 148: Setting the Trust Mode	249
Figure 149: Configuring DSCP to DSCP Internal Mapping	251
Figure 150: Showing DSCP to DSCP Internal Mapping	251
Figure 151: Configuring CoS to DSCP Internal Mapping	253
Figure 152: Showing CoS to DSCP Internal Mapping	253
Figure 153: Configuring DSCP to CoS Egress Mapping	255
Figure 154: Showing DSCP to CoS Egress Mapping	255
Figure 155: Configuring IP Precedence to DSCP Internal Mapping	257
Figure 156: Showing the IP Precedence to DSCP Internal Map	258
Figure 157: Configuring IP Port Number to DSCP Internal Mapping	259
Figure 158: Showing IP Port Number to DSCP Internal Mapping	260
Figure 159: Configuring a Class Map	263
Figure 160: Showing Class Maps	264
Figure 161: Adding Rules to a Class Map	265
Figure 162: Showing the Rules for a Class Map	265
Figure 163: Configuring a Policy Map	273
Figure 164: Showing Policy Maps	273
Figure 165: Adding Rules to a Policy Map	274
Figure 166: Showing the Rules for a Policy Map	274
Figure 167: Attaching a Policy Map to a Port	276
Figure 168: Configuring a Voice VLAN	279
Figure 169: Configuring an OUI Telephony List	280

Figure 170: Showing an OUI Telephony List	280
Figure 171: Configuring Port Settings for a Voice VLAN	282
Figure 172: Configuring the Authentication Sequence	286
Figure 173: Authentication Server Operation	286
Figure 174: Configuring Remote Authentication Server (RADIUS)	289
Figure 175: Configuring Remote Authentication Server (TACACS+)	290
Figure 176: Configuring AAA Server Groups	290
Figure 177: Showing AAA Server Groups	291
Figure 178: Configuring Global Settings for AAA Accounting	293
Figure 179: Configuring AAA Accounting Methods	294
Figure 180: Showing AAA Accounting Methods	294
Figure 181: Configuring AAA Accounting Service for 802.1X Service	295
Figure 182: Configuring AAA Accounting Service for Exec Service	295
Figure 183: Displaying a Summary of Applied AAA Accounting Methods	295
Figure 184: Displaying Statistics for AAA Accounting Sessions	296
Figure 185: Configuring AAA Authorization Methods	297
Figure 186: Showing AAA Authorization Methods	298
Figure 187: Configuring AAA Authorization Methods for Exec Service	298
Figure 188: Displaying the Applied AAA Authorization Method	299
Figure 189: Configuring User Accounts	300
Figure 190: Showing User Accounts	301
Figure 191: Configuring Global Settings for Web Authentication	302
Figure 192: Configuring Interface Settings for Web Authentication	303
Figure 193: Configuring Global Settings for Network Access	307
Figure 194: Configuring Interface Settings for Network Access	309
Figure 195: Configuring Link Detection for Network Access	310
Figure 196: Configuring a MAC Address Filter for Network Access	311
Figure 197: Showing the MAC Address Filter Table for Network Access	311
Figure 198: Showing Addresses Authenticated for Network Access	313
Figure 199: Configuring HTTPS	315
Figure 200: Downloading the Secure-Site Certificate	316
Figure 201: Configuring the SSH Server	320
Figure 202: Generating the SSH Host Key Pair	321
Figure 203: Showing the SSH Host Key Pair	322
Figure 204: Copying the SSH User's Public Key	323

Figure 205: Showing the SSH User's Public Key	324
Figure 206: Showing TCAM Utilization	327
Figure 207: Creating an ACL	328
Figure 208: Showing a List of ACLs	329
Figure 209: Configuring a Standard IPv4 ACL	330
Figure 210: Configuring an Extended IPv4 ACL	333
Figure 211: Configuring a Standard IPv6 ACL	334
Figure 212: Configuring an Extended IPv6 ACL	336
Figure 213: Configuring a MAC ACL	338
Figure 214: Configuring a ARP ACL	340
Figure 215: Binding a Port to an ACL	341
Figure 216: Showing ACL Statistics	342
Figure 217: Creating an IP Address Filter for Management Access	344
Figure 218: Showing IP Addresses Authorized for Management Access	344
Figure 219: Configuring Port Security	346
Figure 220: Configuring Port Authentication	347
Figure 221: Configuring Global Settings for 802.1X Port Authentication	349
Figure 222: Configuring Interface Settings for 802.1X Port Authenticator	353
Figure 223: Showing Statistics for 802.1X Port Authenticator	355
Figure 224: Protecting Against DoS Attacks	356
Figure 225: Configuring Global Settings for DHCP Snooping	360
Figure 226: Configuring DHCP Snooping on a VLAN	361
Figure 227: Configuring the Port Mode for DHCP Snooping	362
Figure 228: Displaying the Binding Table for DHCP Snooping	364
Figure 229: Configuring Global Settings for DHCPv6 Snooping	368
Figure 230: Configuring DHCPv6 Snooping on a VLAN	369
Figure 231: Showing VLANs Enabled for DHCPv6 Snooping	369
Figure 232: Configuring the Trust Sate for DHCPv6 Snooping	370
Figure 233: Displaying the Binding Table for DHCPv6 Snooping	371
Figure 234: Displaying Statistics for DHCPv6 Snooping	372
Figure 235: Setting the Filter Type for IP Source Guard	375
Figure 236: Configuring Static Bindings for IPv4 Source Guard	377
Figure 237: Displaying Static Bindings for IPv4 Source Guard	378
Figure 238: Showing the IPv4 Source Guard Binding Table	379
Figure 239: Setting the Filter Type for IPv6 Source Guard	381

Figure 240: Configuring Static Bindings for IPv6 Source Guard	383
Figure 241: Displaying Static Bindings for IPv6 Source Guard	383
Figure 242: Showing the IPv6 Source Guard Binding Table	385
Figure 243: Configuring Global Settings for ARP Inspection	388
Figure 244: Configuring VLAN Settings for ARP Inspection	389
Figure 245: Configuring Interface Settings for ARP Inspection	390
Figure 246: Displaying Statistics for ARP Inspection	392
Figure 247: Displaying the ARP Inspection Log	393
Figure 248: Configuring Discarding or Forwarding of CDP/PVST Packets	393
Figure 249: Configuring Settings for System Memory Logs	397
Figure 250: Showing Error Messages Logged to System Memory	398
Figure 251: Configuring Settings for Remote Logging of Error Messages	399
Figure 252: Configuring SMTP Alert Messages	401
Figure 253: Configuring LLDP Timing Attributes	403
Figure 254: Configuring LLDP Interface Attributes	407
Figure 255: Configuring the Civic Address for an LLDP Interface	408
Figure 256: Showing the Civic Address for an LLDP Interface	409
Figure 257: Displaying Local Device Information for LLDP (General)	412
Figure 258: Displaying Local Device Information for LLDP (Port)	412
Figure 259: Displaying Local Device Information for LLDP (Port Details)	412
Figure 260: Displaying Remote Device Information for LLDP (Port)	419
Figure 261: Displaying Remote Device Information for LLDP (Port Details)	420
Figure 262: Displaying LLDP Device Statistics (General)	422
Figure 263: Displaying LLDP Device Statistics (Port)	422
Figure 264: Setting the Switch's PoE Budget	424
Figure 265: Setting a Port's PoE Budget	426
Figure 266: Configuring Global Settings for SNMP	429
Figure 267: Setting Community Access Strings	430
Figure 268: Showing Community Access Strings	430
Figure 269: Configuring the Local Engine ID for SNMP	431
Figure 270: Configuring a Remote Engine ID for SNMP	432
Figure 271: Showing Remote Engine IDs for SNMP	433
Figure 272: Creating an SNMP View	434
Figure 273: Showing SNMP Views	434
Figure 274: Adding an OID Subtree to an SNMP View	435

Figure 275: Showing the OID Subtree Configured for SNMP Views	435
Figure 276: Creating an SNMP Group	440
Figure 277: Showing SNMP Groups	440
Figure 278: Configuring Local SNMPv3 Users	442
Figure 279: Showing Local SNMPv3 Users	442
Figure 280: Changing a Local SNMPv3 User Group	443
Figure 281: Configuring Remote SNMPv3 Users	445
Figure 282: Showing Remote SNMPv3 Users	445
Figure 283: Configuring Trap Managers (SNMPv1)	449
Figure 284: Configuring Trap Managers (SNMPv2c)	449
Figure 285: Configuring Trap Managers (SNMPv3)	449
Figure 286: Showing Trap Managers	450
Figure 287: Creating a Trap Host	451
Figure 288: Showing the Notification Filter	451
Figure 289: Creating SNMP Notification Logs	451
Figure 290: Showing SNMP Notification Logs	452
Figure 291: Showing SNMP Statistics	454
Figure 292: Configuring an RMON Alarm	456
Figure 293: Showing Configured RMON Alarms	457
Figure 294: Configuring an RMON Event	458
Figure 295: Showing Configured RMON Events	459
Figure 296: Configuring an RMON History Sample	460
Figure 297: Showing Configured RMON History Samples	461
Figure 298: Showing Collected RMON History Samples	461
Figure 299: Configuring an RMON Statistical Sample	463
Figure 300: Showing Configured RMON Statistical Samples	463
Figure 301: Showing Collected RMON Statistical Samples	464
Figure 302: Configuring a Switch Cluster	466
Figure 303: Configuring a Cluster Members	467
Figure 304: Showing Cluster Members	467
Figure 305: Showing Cluster Candidates	467
Figure 306: Managing a Cluster Member	468
Figure 307: Setting the Name of a Time Range	470
Figure 308: Showing a List of Time Ranges	470
Figure 309: Add a Rule to a Time Range	471

Figure 310: Showing the Rules Configured for a Time Range	471
Figure 311: ERPS Ring Components	473
Figure 312: Ring Interconnection Architecture (Multi-ring/Ladder Network)	474
Figure 313: Setting ERPS Global Status	476
Figure 314: Sub-ring with Virtual Channel	486
Figure 315: Sub-ring without Virtual Channel	486
Figure 316: Non-ERPS Device Protection	487
Figure 317: Creating an ERPS Ring	490
Figure 318: Creating an ERPS Ring	491
Figure 319: Showing Configured ERPS Rings	492
Figure 320: Blocking an ERPS Ring Port	496
Figure 321: Enabling OAM for Local Ports	499
Figure 322: Displaying Statistics for OAM Messages	500
Figure 323: Displaying the OAM Event Log	501
Figure 324: Displaying Status of Remote Interfaces	502
Figure 325: Running a Remote Loop Back Test	504
Figure 326: Displaying the Results of Remote Loop Back Testing	505
Figure 327: Single CFM Maintenance Domain	506
Figure 328: Multiple CFM Maintenance Domains	507
Figure 329: Configuring Global Settings for CFM	511
Figure 330: Configuring Interfaces for CFM	512
Figure 331: Configuring Maintenance Domains	516
Figure 332: Showing Maintenance Domains	516
Figure 333: Configuring Detailed Settings for Maintenance Domains	517
Figure 334: Creating Maintenance Associations	520
Figure 335: Showing Maintenance Associations	520
Figure 336: Configuring Detailed Settings for Maintenance Associations	521
Figure 337: Configuring Maintenance End Points	522
Figure 338: Showing Maintenance End Points	523
Figure 339: Configuring Remote Maintenance End Points	524
Figure 340: Showing Remote Maintenance End Points	525
Figure 341: Transmitting Link Trace Messages	527
Figure 342: Transmitting Loopback Messages	528
Figure 343: Transmitting Delay-Measure Messages	530
Figure 344: Showing Information on Local MEPs	531

Figure 345: Showing Detailed Information on Local MEPs	533
Figure 346: Showing Information on Local MIPs	534
Figure 347: Showing Information on Remote MEPs	535
Figure 348: Showing Detailed Information on Remote MEPs	537
Figure 349: Showing the Link Trace Cache	539
Figure 350: Showing Settings for the Fault Notification Generator	540
Figure 351: Showing Continuity Check Errors	541
Figure 352: Enabling OAM for Local Ports	544
Figure 353: Displaying Statistics for OAM Messages	545
Figure 354: Displaying the OAM Event Log	546
Figure 355: Displaying Status of Remote Interfaces	547
Figure 356: Running a Remote Loop Back Test	549
Figure 357: Displaying the Results of Remote Loop Back Testing	550
Figure 358: Configuring UDLD Protocol Intervals	552
Figure 359: Configuring UDLD Interface Settings	554
Figure 360: Displaying UDLD Neighbor Information	555
Figure 361: Configuring Global Settings for LBD	557
Figure 362: Configuring Interface Settings for LBD	558
Figure 363: Multicast Filtering Concept	560
Figure 364: Configuring General Settings for IGMP Snooping	566
Figure 365: Configuring a Static Interface for a Multicast Router	567
Figure 366: Showing Static Interfaces Attached a Multicast Router	568
Figure 367: Showing Current Interfaces Attached a Multicast Router	568
Figure 368: Assigning an Interface to a Multicast Service	569
Figure 369: Showing Static Interfaces Assigned to a Multicast Service	570
Figure 370: Configuring IGMP Snooping on a VLAN	575
Figure 371: Showing Interface Settings for IGMP Snooping	576
Figure 372: Dropping IGMP Query or Multicast Data Packets	577
Figure 373: Showing Multicast Groups Learned by IGMP Snooping	578
Figure 374: Displaying IGMP Snooping Statistics – Query	580
Figure 375: Displaying IGMP Snooping Statistics – VLAN	581
Figure 376: Displaying IGMP Snooping Statistics – Port	582
Figure 377: Enabling IGMP Filtering and Throttling	583
Figure 378: Creating an IGMP Filtering Profile	584
Figure 379: Showing the IGMP Filtering Profiles Created	585

Figure 380: Adding Multicast Groups to an IGMP Filtering Profile	585
Figure 381: Showing the Groups Assigned to an IGMP Filtering Profile	586
Figure 382: Configuring IGMP Filtering and Throttling Interface Settings	587
Figure 383: Configuring General Settings for MLD Snooping	589
Figure 384: Configuring Immediate Leave for MLD Snooping	590
Figure 385: Configuring a Static Interface for an IPv6 Multicast Router	591
Figure 386: Showing Static Interfaces Attached an IPv6 Multicast Router	591
Figure 387: Showing Current Interfaces Attached an IPv6 Multicast Router	592
Figure 388: Assigning an Interface to an IPv6 Multicast Service	593
Figure 389: Showing Static Interfaces Assigned to an IPv6 Multicast Service	593
Figure 390: Showing Current Interfaces Assigned to an IPv6 Multicast Service	594
Figure 391: Showing IPv6 Multicast Services and Corresponding Sources	595
Figure 392: MVR Concept	596
Figure 393: Configuring Global Settings for MVR	599
Figure 394: Configuring Domain Settings for MVR	600
Figure 395: Configuring an MVR Group Address Profile	601
Figure 396: Displaying MVR Group Address Profiles	602
Figure 397: Assigning an MVR Group Address Profile to a Domain	602
Figure 398: Showing the MVR Group Address Profiles Assigned to a Domain	603
Figure 399: Configuring Interface Settings for MVR	605
Figure 400: Assigning Static MVR Groups to an Interface	607
Figure 401: Showing the Static MVR Groups Assigned to a Port	607
Figure 402: Displaying MVR Receiver Groups	608
Figure 403: Displaying MVR Statistics – Query	610
Figure 404: Displaying MVR Statistics – VLAN	611
Figure 405: Displaying MVR Statistics – Port	612
Figure 406: Configuring Global Settings for MVR6	614
Figure 407: Configuring Domain Settings for MVR6	616
Figure 408: Configuring an MVR6 Group Address Profile	617
Figure 409: Displaying MVR6 Group Address Profiles	618
Figure 410: Assigning an MVR6 Group Address Profile to a Domain	618
Figure 411: Showing MVR6 Group Address Profiles Assigned to a Domain	619
Figure 412: Configuring Interface Settings for MVR6	621
Figure 413: Assigning Static MVR6 Groups to a Port	622
Figure 414: Showing the Static MVR6 Groups Assigned to a Port	623

Figure 415: Displaying MVR6 Receiver Groups	624
Figure 416: Displaying MVR6 Statistics – Query	626
Figure 417: Displaying MVR6 Statistics – VLAN	627
Figure 418: Displaying MVR6 Statistics – Port	628
Figure 419: Pinging a Network Device	630
Figure 420: Tracing the Route to a Network Device	632
Figure 421: Proxy ARP	633
Figure 422: Configuring General Settings for ARP	634
Figure 423: Configuring Static ARP Entries	635
Figure 424: Displaying Static ARP Entries	636
Figure 425: Displaying ARP Entries	636
Figure 426: Displaying ARP Statistics	637
Figure 427: Configuring a Static IPv4 Address	641
Figure 428: Configuring a Dynamic IPv4 Address	642
Figure 429: Showing the Configured IPv4 Address for an Interface	643
Figure 430: Configuring the IPv6 Default Gateway	644
Figure 431: Configuring General Settings for an IPv6 Interface	648
Figure 432: Configuring RA Guard for an IPv6 Interface	649
Figure 433: Configuring an IPv6 Address	651
Figure 434: Showing Configured IPv6 Addresses	653
Figure 435: Showing IPv6 Neighbors	654
Figure 436: Showing IPv6 Statistics (IPv6)	658
Figure 437: Showing IPv6 Statistics (ICMPv6)	659
Figure 438: Showing IPv6 Statistics (UDP)	659
Figure 439: Showing Reported MTU Values	660
Figure 440: Virtual Interfaces and Layer 3 Routing	662
Figure 441: Configuring Static Routes	666
Figure 442: Displaying Static Routes	666
Figure 443: Displaying the Routing Table	667
Figure 444: Configuring General Settings for DNS	670
Figure 445: Configuring a List of Domain Names for DNS	671
Figure 446: Showing the List of Domain Names for DNS	671
Figure 447: Configuring a List of Name Servers for DNS	672
Figure 448: Showing the List of Name Servers for DNS	673
Figure 449: Configuring Static Entries in the DNS Table	674

Figure 450: Showing Static Entries in the DNS Table	674
Figure 451: Showing Entries in the DNS Cache	675
Figure 452: Specifying A DHCP Client Identifier	677
Figure 453: Layer 3 DHCP Relay Service	677
Figure 454: Configuring DHCP Relay Service	678
Figure 455: Enabling Dynamic Provisioning via DHCP	679
Figure 456: Configuring Global Settings for PPPoE Intermediate Agent	680
Figure 457: Configuring Interface Settings for PPPoE Intermediate Agent	682
Figure 458: Showing PPPoE Intermediate Agent Statistics	683

Tables

Table 1: Key Features	35
Table 2: System Defaults	42
Table 3: Web Page Configuration Buttons	50
Table 4: Switch Main Menu	52
Table 5: Predefined Summer-Time Parameters	94
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	217
Table 12: Default STA Path Costs	217
Table 13: IEEE 802.1p Egress Queue Priority Mapping	245
Table 14: CoS Priority Levels	246
Table 15: Mapping Internal Per-hop Behavior to Hardware Queues	246
Table 16: Default Mapping of DSCP Values to Internal PHB/Drop Values	250
Table 17: Default Mapping of CoS/CFI to Internal PHB/Drop Precedence	252
Table 18: Mapping Internal PHB/Drop Precedence to CoS/CFI Values	254
Table 19: Mapping IP Precedence	256
Table 20: Default Mapping of IP Precedence to Internal PHB/Drop Values	257
Table 21: Dynamic QoS Profiles	305
Table 22: HTTPS System Support	314
Table 23: 802.1X Statistics	353
Table 24: ARP Inspection Statistics	391
Table 25: ARP Inspection Log	392
Table 26: Logging Levels	396
Table 27: LLDP MED Location CA Types	407
Table 28: Chassis ID Subtype	409
Table 29: System Capabilities	410

Table 30: Port ID Subtype	411
Table 31: Remote Port Auto-Negotiation Advertised Capability	414
Table 32: Maximum Number of Ports Providing Simultaneous Power	425
Table 33: SNMPv3 Security Models and Levels	427
Table 34: Supported Notification Messages	436
Table 35: ERPS Request/State Priority	493
Table 36: OAM Operation State	497
Table 37: Remote Loopback Status	503
Table 38: Remote MEP Priority Levels	514
Table 39: MEP Defect Descriptions	514
Table 40: OAM Operation State	542
Table 41: Remote Loopback Status	548
Table 42: Address Resolution Protocol	632
Table 43: ARP Statistics	636
Table 44: Show IPv6 Neighbors - display description	653
Table 45: Show IPv6 Statistics - display description	655
Table 46: Show MTU - display description	660
Table 47: Options 60, 66 and 67 Statements	676
Table 48: Options 55 and 124 Statements	676
Table 49: Troubleshooting Chart	693

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 35](#)

Introduction

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

Key Features

Table 1: Key Features

Feature	Description
Configuration Backup and Restore	Using management station or FTP/TFTP server
Authentication	Console, Telnet, web – user name/password, RADIUS, TACACS+ Port – IEEE 802.1X, MAC address filtering SNMP v1/2c - Community strings SNMP version 3 – MD5 or SHA password Telnet – SSH Web – HTTPS
General Security Measures	AAA ARP inspection DHCP Snooping (with Option 82 relay information) IP Source Guard PPPoE Intermediate Agent Private VLANs Port Authentication – IEEE 802.1X Port Security – MAC address filtering
Access Control Lists	Supports up to 256 ACLs, up to 1K rules per ACL
DHCP	Client, Relay
DHCPv6	Client
DNS	Client and Proxy service
Port Configuration	Speed, duplex mode and flow control
Port Trunking	Supports up to 16/27 trunks – static or dynamic trunking (LACP)
Port Mirroring	1 session, one or more source ports to one analysis port
Congestion Control	Rate Limiting Throttling for broadcast, multicast, unknown unicast storms

Table 1: Key Features (Continued)

Feature	Description
Address Table	16K MAC addresses in forwarding table, 1K static MAC addresses; 1K entries in ARP cache; 2K entries in ipv6 neighbor cache; 1K L2 IPv4 multicast groups
IP Version 4 and 6	Supports IPv4 and IPv6 addressing, and management
IEEE 802.1D Bridge	Supports dynamic data switching and addresses learning
Store-and-Forward Switching	Supported to ensure wire-speed switching while eliminating bad frames
Spanning Tree Algorithm	Supports standard STP, Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Trees (MSTP)
Virtual LANs	Up to 4094 using IEEE 802.1Q, port-based, protocol-based, private VLANs, voice VLANs, and QinQ tunnel
Traffic Prioritization	Default port priority, traffic class map, queue scheduling, IP Precedence, or Differentiated Services Code Point (DSCP), and TCP/UDP Port
Qualify of Service	Supports Differentiated Services (DiffServ)
Link Layer Discovery Protocol	Used to discover basic information about neighboring devices
Switch Clustering	Supports up to 16 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)
ARP	Static and dynamic address configuration, proxy ARP
Multicast Filtering	Supports IGMP snooping and query for Layer 2, MLD snooping and query, IGMP for Layer 3, and Multicast VLAN Registration
Remote Device Management	Supports Ethernet OAM functions for attached CPEs (IEEE 802.3ah, ITU-T Y.1731)

Description of Software Features

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

Some of the management features are briefly described below.

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

Authentication This switch authenticates management access via the console port, Telnet, or a web browser. User names and passwords can be configured locally or can be verified via a remote authentication server (i.e., RADIUS or TACACS+). Port-based authentication is also supported via the IEEE 802.1X protocol. This protocol uses Extensible Authentication Protocol over LANs (EAPOL) to request user credentials from the 802.1X client, and then uses the EAP between the switch and the authentication server to verify the client's right to access the network via an authentication server (i.e., RADIUS or TACACS+ server).

Other authentication options include HTTPS for secure management access via the web, SSH for secure management access over a Telnet-equivalent connection, SNMP Version 3, IP address filtering for SNMP/Telnet/web management access. MAC address filtering and IP source guard also provide authenticated port access. While DHCP snooping is provided to prevent malicious attacks from insecure ports. While PPPoE Intermediate Agent supports authentication of a client for a service provider.

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

DHCP Relay Since DHCP uses a broadcast mechanism, a DHCP server and its client must physically reside on the same subnet. Since it is not practical to have a DHCP server on every subnet, DHCP Relay is supported to allow dynamic configuration of local clients from a DHCP server located in a different network. DHCP Relay Option 82 is also provided to control the processing of Option 82 information in DHCP request packets relayed by this device.

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

Rate Limiting This feature controls the maximum rate for traffic transmitted or received on an interface. Rate limiting is configured on interfaces at the edge of a network to limit traffic into or out of the network. Packets that exceed the acceptable amount of traffic are dropped.

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

Port Trunking Ports can be combined into an aggregate connection. Trunks can be manually set up or dynamically configured using Link Aggregation Control Protocol (LACP – IEEE 802.3-2005). The additional ports dramatically increase the throughput across any connection, and provide redundancy by taking over the load if a port in the trunk should fail. The switch supports up to 16/27 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 3 Mbits for frame buffering. This buffer can queue packets awaiting transmission on congested networks.

Spanning Tree Algorithm

The switch supports these spanning tree protocols:

- ◆ Spanning Tree Protocol (STP, IEEE 802.1D) – This protocol provides loop detection. When there are multiple physical paths between segments, this protocol will choose a single path and disable all others to ensure that only one route exists between any two stations on the network. This prevents the creation of network loops. However, if the chosen path should fail for any reason, an alternate path will be activated to maintain the connection.
- ◆ Rapid Spanning Tree Protocol (RSTP, IEEE 802.1w) – This protocol reduces the convergence time for network topology changes to about 3 to 5 seconds, compared to 30 seconds or more for the older IEEE 802.1D STP standard. It is intended as a complete replacement for STP, but can still interoperate with switches running the older standard by automatically reconfiguring ports to STP-compliant mode if they detect STP protocol messages from attached devices.
- ◆ Multiple Spanning Tree Protocol (MSTP, IEEE 802.1s) – This protocol is a direct extension of RSTP. It can provide an independent spanning tree for different VLANs. It simplifies network management, provides for even faster convergence than RSTP by limiting the size of each region, and prevents VLAN members from being segmented from the rest of the group (as sometimes occurs with IEEE 802.1D STP).

Connectivity Fault Management

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

Virtual LANs

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

- ◆ Eliminate broadcast storms which severely degrade performance in a flat network.
- ◆ Simplify network management for node changes/moves by remotely configuring VLAN membership for any port, rather than having to manually change the network connection.

- ◆ Provide data security by restricting all traffic to the originating VLAN, except where a connection is explicitly defined via the switch's routing service.
- ◆ Use private VLANs to restrict traffic to pass only between data ports and the uplink ports, thereby isolating adjacent ports within the same VLAN, and allowing you to limit the total number of VLANs that need to be configured.
- ◆ Use protocol VLANs to restrict traffic to specified interfaces based on protocol type.

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

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

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

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

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

- Address Resolution Protocol** The switch uses ARP to convert between IP addresses and MAC (hardware) addresses. This switch supports conventional ARP, which locates the MAC address corresponding to a given IP address.
- Operation, Administration, and Maintenance** The switch provides OAM remote management tools required to monitor and maintain the links to subscriber CPEs (Customer Premise Equipment). This section describes functions including enabling OAM for selected ports, loopback testing, and displaying remote device information.
- Multicast Filtering** Specific multicast traffic can be assigned to its own VLAN to ensure that it does not interfere with normal network traffic and to guarantee real-time delivery by setting the required priority level for the designated VLAN. The switch uses IGMP Snooping and Query for IPv4, MLD Snooping and Query for IPv6, and IGMP at Layer 3 to manage multicast group registration. It also supports Multicast VLAN Registration (MVR for IPv4 and MVR6 for IPv6) which allows common multicast traffic, such as television channels, to be transmitted across a single network-wide multicast VLAN shared by hosts residing in other standard or private VLAN groups, while preserving security and data isolation for normal traffic.
- Link Layer Discovery Protocol** LLDP is used to discover basic information about neighboring devices within the local broadcast domain. LLDP is a Layer 2 protocol that advertises information about the sending device and collects information gathered from neighboring network nodes it discovers.
- Advertised information is represented in Type Length Value (TLV) format according to the IEEE 802.1ab standard, and can include details such as device identification, capabilities and configuration settings. Media Endpoint Discovery (LLDP-MED) is an extension of LLDP intended for managing endpoint devices such as Voice over IP phones and network switches. The LLDP-MED TLVs advertise information such as network policy, power, inventory, and device location details. The LLDP and LLDP-MED information can be used by SNMP applications to simplify troubleshooting, enhance network management, and maintain an accurate network topology.

System Defaults

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

The following table lists some of the basic system defaults.

Table 2: System Defaults

Function	Parameter	Default
Console Port Connection	Baud Rate	115200 bps
	Data bits	8
	Stop bits	1
	Parity	none
	Local Console Timeout	0 (disabled)
Authentication and Security Measures	Privileged Exec Level	Username "admin" Password "admin"
	Normal Exec Level	Username "guest" Password "guest"
	Enable Privileged Exec from Normal Exec Level	Password "super"
	RADIUS Authentication	Disabled
	TACACS+ Authentication	Disabled
	802.1X Port Authentication	Disabled
	MAC Authentication	Disabled
	PPPoE Intermediate Agent	Disabled
	HTTPS	Enabled
	SSH	Disabled
	Port Security	Disabled
	IP Filtering	Disabled
	DHCP Snooping	Disabled
	Web Management	HTTP Server
Web Management	HTTP Port Number	80
	HTTP Secure Server	Enabled
	HTTP Secure Server Port	443

Table 2: System Defaults (Continued)

Function	Parameter	Default
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
Congestion Control	Rate Limiting	Disabled
	Storm Control	Broadcast: Enabled (500 packets/sec) Multicast: Disabled Unknown Unicast: Disabled
Address Table	Aging Time	300 seconds
Spanning Tree Algorithm	Status	Enabled, RSTP (Defaults: RSTP standard)
	Edge Ports	Disabled
LLDP	Status	Enabled
ERPS	Status	Disabled
CFM	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

Table 2: System Defaults (Continued)

Function	Parameter	Default
Traffic Prioritization	Ingress Port Priority	0
	Queue Mode	WRR
	Queue Weight	Queue: 0 1 2 3 4 5 6 7 Weight: 1 2 4 6 8 10 12 14
	Class of Service	Enabled
	IP Precedence Priority	Disabled
	IP DSCP Priority	Disabled
	IP Port Priority	Disabled
IP Settings	Management. VLAN	VLAN 1
	IP Address	DHCP assigned
	Subnet Mask	255.255.255.0
	Default Gateway	0.0.0.0
	DHCP	Client: Enabled
	DNS	Client/Proxy service: Disabled
	BOOTP	Disabled
	ARP	Enabled Cache Timeout: 20 minutes Proxy: Disabled
Multicast Filtering	IGMP Snooping (Layer 2)	Snooping: Enabled Querier: Disabled
	MLD Snooping (Layer 2 IPv6)	Snooping: Enabled Querier: Disabled
	Multicast VLAN Registration	Disabled
	IGMP Proxy Reporting	Enabled
	IGMP (Layer 3)	Disabled
	IGMP Proxy (Layer 3)	Disabled
System Log	Status	Enabled
	Messages Logged to RAM	Levels 0-7 (all)
	Messages Logged to Flash	Levels 0-3
SMTP Email Alerts	Event Handler	Enabled (but no server defined)
SNTP	Clock Synchronization	Disabled
Switch Clustering	Status	Disabled
	Commander	Disabled

Section II

Web Configuration

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

This section includes these chapters:

- ◆ ["Using the Web Interface" on page 47](#)
- ◆ ["Basic Management Tasks" on page 71](#)
- ◆ ["Interface Configuration" on page 105](#)
- ◆ ["VLAN Configuration" on page 159](#)
- ◆ ["Address Table Settings" on page 197](#)
- ◆ ["Spanning Tree Algorithm" on page 205](#)
- ◆ ["Congestion Control" on page 231](#)
- ◆ ["Class of Service" on page 241](#)
- ◆ ["Quality of Service" on page 261](#)
- ◆ ["VoIP Traffic Configuration" on page 277](#)
- ◆ ["Security Measures" on page 283](#)
- ◆ ["Basic Administration Protocols" on page 395](#)
- ◆ ["Multicast Filtering" on page 559](#)
- ◆ ["IP Configuration" on page 639](#)
- ◆ ["General IP Routing" on page 661](#)
- ◆ ["IP Services" on page 669](#)
- ◆ ["Basic IP Functions" on page 629](#)

- ◆ ["Configuring Router Redundancy" on page 579](#)

2

Using the Web Interface

This switch provides an embedded HTTP web agent. Using a web browser you can configure the switch and view statistics to monitor network activity. The web agent can be accessed by any computer on the network using a standard web browser (Internet Explorer 11, Mozilla Firefox 43, or Google Chrome 47, or more recent versions).



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

Connecting to the Web Interface

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

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



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

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

Note: If the path between your management station and this switch does not pass through any device that uses the Spanning Tree Algorithm, then you can set the switch port attached to your management station to fast forwarding (i.e., enable Admin Edge Port) to improve the switch’s response time to management commands issued through the web interface. See [“Configuring Interface Settings for STA” on page 216](#).


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 299](#) 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. Switch Information, CPU Utilization, Switch Events, Memory Utilization, Recent 5 Event Information, Port Utilization, Dynamic Address Count, and LLDP Remote Device Port List are displayed on the right side. The main menu links are used to navigate to other menus, and display configuration parameters and statistics.

Figure 1: Dashboard






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

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

Figure 2: Home Page



NOTE: This manual covers the ECS4120-28T/52T Gigabit Ethernet switches, the ECS4120-28F Gigabit Ethernet fiber switch, and the ECS4120-28P Gigabit Ethernet PoE switch. Other than the difference in port types, and support for PoE (ECS4120-28P), there are no significant differences. Therefore most of the screen display examples are based on the ECS4120-28T. The panel graphics for the various switch types are shown on the following page.

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

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

Table 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.”

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 Ports (i.e., up or down), Duplex State (i.e., half or full duplex), or Flow Control (i.e., with or without flow control).

Figure 3: Front Panel Indicators

ECS4120-28T



ECS4120-28P



ECS4120-28F



ECS4120-52T



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	Display switch Information, CPU utilization, switch events, memory utilization, recent 5 event information, port utilization, dynamic address count, and LLDP remote device port list	48
System		
General	Provides basic system description, including contact information	72
Switch	Shows the number of ports, hardware version, power status, and firmware version numbers	73
Capability	Enables support for jumbo frames; shows the bridge extension parameters	74 , 75
File		77
Copy	Allows the transfer and copying files	77
Set Startup	Sets the startup file	80
Show	Shows the files stored in flash memory; allows deletion of files	81
Time		85
Configure General		
Manual	Manually sets the current time	86
SNTP	Configures SNTP polling interval	87
NTP	Configures NTP authentication parameters	87
Configure Time Server	Configures a list of NTP or SNTP servers	88
Configure SNTP Server	Sets the IP address for SNTP time servers	88
Add NTP Server	Adds NTP time server and index of authentication key	89
Show NTP Server	Shows list of configured NTP time servers	89
Add NTP Authentication Key	Adds key index and corresponding MD5 key	91
Show NTP Authentication Key	Shows list of configured authentication keys	91
Configure Time Zone	Sets the local time zone for the system clock	92
Configure Summer Time	Configures summer time settings	93
Console	Sets console port connection parameters	95
Telnet	Sets Telnet connection parameters	97
CPU Utilization	Displays information on CPU utilization	98
CPU Guard	Sets the CPU utilization watermark and threshold	99
Memory Status	Shows memory utilization parameters	100
Reset	Restarts the switch immediately, at a specified time, after a specified delay, or at a periodic interval	101

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Interface		105
Port		106
General		
Configure by Port List	Configures connection settings per port	106
Configure by Port Range	Configures connection settings for a range of ports	109
Show Information	Displays port connection status	110
Statistics	Shows Interface, Etherlike, and RMON port statistics	111
Charts	Shows Interface, Etherlike, RMON, and all port statistics	111
History	Shows statistical history for specified interfaces	116
Transceiver	Configures thresholds for alarm and warning messages for optical transceivers which support DDM	120
Cable Test	Performs cable diagnostics for selected port to diagnose any cable faults (short, open etc.) and report the cable length	123
Trunk		
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		129
Configure		129
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
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

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
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
Add	Assign the downlink and uplink ports to use in a segmented group	154
Show	Shows the assigned ports and direction (uplink/downlink)	154
VLAN Trunking	Allows unknown VLAN groups to pass through the specified interface	156
VLAN	Virtual LAN	159
Static		
Add	Creates VLAN groups	163
Show	Displays configured VLAN groups	163
Modify	Configures group name and administrative status	163
Edit Member by VLAN	Specifies VLAN attributes per VLAN	165
Edit Member by Interface	Specifies VLAN attributes per interface	165
Edit Member by Interface Range	Specifies VLAN attributes per interface range	165

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Dynamic		
Configure General	Enables GVRP VLAN registration protocol globally	170
Configure Interface	Configures GVRP status and timers per interface	170
Show Dynamic VLAN		170
Show VLAN	Shows the VLANs this switch has joined through GVRP	170
Show VLAN Member	Shows the interfaces assigned to a VLAN through GVRP	170
Tunnel	IEEE 802.1Q (QinQ) Tunneling	173
Configure Global	Sets tunnel mode for the switch	177
Configure Service	Sets a CVLAN to SPVLAN mapping entry	178
Configure Interface	Sets the tunnel mode for any participating interface	180
L2PT	Layer 2 Protocol Tunneling	181
Configure Global	Configures the destination MAC address for L2PT	183
Configure Interface	Enables L2PT on selected interfaces	184
Protocol		185
Configure Protocol		186
Add	Creates a protocol group, specifying supported protocols	186
Show	Shows configured protocol groups	186
Configure Interface		187
Add	Maps a protocol group to a VLAN	187
Show	Shows the protocol groups mapped to each VLAN	187
IP Subnet		189
Add	Maps IP subnet traffic to a VLAN	189
Show	Shows IP subnet to VLAN mapping	189
MAC-Based		191
Add	Maps traffic with specified source MAC address to a VLAN	191
Show	Shows source MAC address to VLAN mapping	191
Translation		193
Add	Maps VLAN IDs between the customer and service provider	193
Show	Displays the configuration settings for VLAN translation	193
MAC Address		197
Dynamic		
Configure Aging	Sets timeout for dynamically learned entries	199
Show Dynamic MAC	Displays dynamic entries in the address table	197

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Clear Dynamic MAC	Removes any learned entries from the forwarding database and clears the transmit and receive counts for any static or system configured entries	198
Learning Status	Enables MAC address learning on selected interfaces	200
Static		201
Add	Configures static entries in the address table	201
Show	Displays static entries in the address table	201
MAC Notification		203
Configure Global	Issues a trap when a dynamic MAC address is added or removed	203
Configure Interface	Enables MAC authentication traps on the current interface	203
Spanning Tree		205
Loopback Detection	Configures Loopback Detection parameters	208
STA	Spanning Tree Algorithm	
Configure Global		
Configure	Configures global bridge settings for STP, RSTP and MSTP	209
Show Information	Displays STA values used for the bridge	215
Configure Interface		
Configure	Configures interface settings for STA	216
Show Inform at on	Displays interface settings for STA	220
MSTP	Multiple Spanning Tree Algorithm	223
Configure Global		223
Add	Configures initial VLAN and priority for an MST instance	223
Show	Configures global settings for an MST instance	223
Modify	Configures the priority or an MST instance	223
Add Member	Adds VLAN members for an MST instance	223
Show Member	Adds or deletes VLAN members for an MST instance	223
Show Information	Displays MSTP values used for the bridge	
Configure Interface		227
Configure	Configures interface settings for an MST instance	227
Show Information	Displays interface settings for an MST instance	227
Traffic		
Rate Limit	Sets the input and output rate limits for a port	231
Storm Control	Sets the broadcast storm threshold for each interface	232

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
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	234
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	236
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	237
Priority		
Default Priority	Sets the default priority for each port or trunk	241
Queue	Sets queue mode for the switch; sets the service weight for each queue that will use a weighted or hybrid mode	242
Trust Mode	Selects DSCP or CoS priority processing	248
DSCP to DSCP		249
Add	Maps DSCP values in incoming packets to per-hop behavior and drop precedence values for internal priority processing	249
Show	Shows the DSCP to DSCP mapping list	249
CoS to DSCP		252
Configure	Maps CoS/CFI values in incoming packets to per-hop behavior and drop precedence values for priority processing	252
Show	Shows the CoS to DSCP mapping list	252
DSCP to CoS		254
Add	Maps internal per-hop behavior and drop precedence value pairs to CoS values used in tagged egress packets on a Layer 2 interface	254
Show	Shows the DSCP to CoS mapping list	254
IP Precedence to DSCP		256
Add	Maps IP precedence values in incoming packets to per-hop behavior and drop precedence values for priority processing	256
Show	Shows the IP Precedence to DSCP mapping list	256
IP Port to DSCP		258
Add	Sets TCP/UDP port priority, defining the socket number and associated per-hop behavior and drop precedence	258
Show	Shows the IP Port to DSCP mapping list	258
PHB to Queue		245
Configure	Maps internal per-hop behavior values to hardware queues	245
Show	Shows the PHB to Queue mapping list	245
DiffServ		261
Configure Class		262
Add	Creates a class map for a type of traffic	262

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show	Shows configured class maps	262
Modify	Modifies the name of a class map	262
Add Rule	Configures the criteria used to classify ingress traffic	262
Show Rule	Shows the traffic classification rules for a class map	262
Configure Policy		265
Add	Creates a policy map to apply to multiple interfaces	265
Show	Shows configured policy maps	265
Modify	Modifies the name of a policy map	265
Add Rule	Sets the boundary parameters used for monitoring inbound traffic, and the action to take for conforming and non-conforming traffic	265
Show Rule	Shows the rules used to enforce bandwidth policing for a policy map	265
Configure Interface	Applies a policy map to an ingress port	275
VoIP	Voice over IP	277
Configure Global	Configures auto-detection of VoIP traffic, sets the Voice VLAN, and VLAN aging time	278
Configure OUI		279
Add	Maps the OUI in the source MAC address of ingress packets to the VoIP device manufacturer	279
Show	Shows the OUI telephony list	279
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	280
Security		283
AAA	Authentication, Authorization and Accounting	284
System Authentication	Configures authentication sequence – local, RADIUS, and TACACS	285
Server		286
Configure Server	Configures RADIUS and TACACS server message exchange settings	286
Configure Group		286
Add	Specifies a group of authentication servers and sets the priority sequence	286
Show	Shows the authentication server groups and priority sequence	286
Accounting	Enables accounting of requested services for billing or security purposes	291
Configure Global	Specifies the interval at which the local accounting service updates information to the accounting server	291
Configure Method		291
Add	Configures accounting for various service types	291
Show	Shows the accounting settings used for various service types	291

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Configure Service	Sets the accounting method applied to specific interfaces for 802.1X, CLI command privilege levels for the console port, and for Telnet	291
Show Information		291
Summary	Shows the configured accounting methods, and the methods applied to specific interfaces	291
Statistics	Shows basic accounting information recorded for user sessions	291
Authorization	Enables authorization of requested services	296
Configure Method		296
Add	Configures authorization for various service types	296
Show	Shows the authorization settings used for various service types	296
Configure Service	Sets the authorization method applied used for the console port, and for Telnet	296
Show Information	Shows the configured authorization methods, and the methods applied to specific interfaces	296
User Accounts		299
Add	Configures user names, passwords, and access levels	299
Show	Shows authorized users	299
Modify	Modifies user attributes	299
Web Authentication	Allows authentication and access to the network when 802.1X or Network Access authentication are infeasible or impractical	301
Configure Global	Configures general protocol settings	301
Configure Interface	Enables Web Authentication for individual ports	302
Network Access	MAC address-based network access authentication	304
Configure Global	Enables aging for authenticated MAC addresses, and sets the time period after which a connected MAC address must be reauthenticated	306
Configure Interface		307
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	307
Link Detection	Configures detection of changes in link status, and the response (i.e., send trap or shut down port)	309
Configure MAC Filter		310
Add	Specifies MAC addresses exempt from authentication	310
Show	Shows the list of exempt MAC addresses	310
Show Information	Shows the authenticated MAC address list	312
HTTPS	Secure HTTP	313
Configure Global	Enables HTTPSs, and specifies the UDP port to use	313
Copy Certificate	Replaces the default secure-site certificate	315

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
SSH	Secure Shell	316
Configure Global	Configures SSH server settings	319
Configure Host Key		320
Generate	Generates the host key pair (public and private)	320
Show	Displays RSA and DSA host keys; deletes host keys	320
Configure User Key		322
Copy	Imports user public keys from TFTP server	322
Show	Displays RSA and DSA user keys; deletes user keys	322
ACL	Access Control Lists	324
Configure ACL		
Show TCAM	Shows utilization parameters for TCAM	325
Add	Adds an ACL based on IP or MAC address filtering	327
Show	Shows the name and type of configured ACLs	327
Add Rule	Configures packet filtering based on IP or MAC addresses and other packet attributes	329, 330, 333, 334, 336, 338
Show Rule	Shows the rules specified for an ACL	
Configure Interface		340
Configure	Binds a port to the specified ACL and time range	340
Show Hardware Counter	Shows statistics for ACL hardware counters	341
IP Filter		342
Add	Sets IP addresses of clients allowed management access via the web, SNMP, and Telnet	342
Show	Shows the addresses to be allowed management access	342
Port Security	Configures per port security, including status, response for security breach, and maximum allowed MAC addresses	344
Port Authentication	IEEE 802.1X	347
Configure Global	Enables authentication and EAPOL pass-through	348
Configure Interface	Sets authentication parameters for individual ports	349
Show Statistics	Displays protocol statistics for the selected port	353
DoS Protection	Protects against Denial-of-Service attacks	355
DHCP Snooping		356
Configure Global	Enables DHCPv4 snooping globally, MAC-address verification, information option; and sets the information policy	359
Configure VLAN	Enables DHCPv4 snooping on a VLAN	360
Configure Interface	Sets the trust mode for an interface	361
Show Information	Displays the DHCPv4 Snooping binding information	363

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
DHCP Snooping6		364
Configure Global	Enables DHCPv6 snooping globally, information option; and sets the information policy	366
Configure VLAN	Enables DHCPv6 snooping on a VLAN	368
Configure Interface	Sets the trust mode for an interface	369
Show Information		
Binding	Displays the DHCPv6 Snooping binding information	371
Statistics	Displays information on client, server, and relay packets	372
IP Source Guard	Filters IPv4 traffic based on static entries in the IP Source Guard table, or dynamic entries in the DHCPv4 Snooping table	373
General	Enables IP source guard, selects filter type per port, and sets maximum binding entries	373
Static Binding		375
Configure ACL Table		375
Add	Adds a static addresses to the source-guard binding table	375
Show	Shows static addresses in the source-guard binding table	375
Configure MAC Table		375
Add	Adds a static addresses to the source-guard binding table	375
Show	Shows static addresses in the source-guard binding table	375
Dynamic Binding	Displays the source-guard binding table for a selected interface	378
IPv6 Source Guard	Filters IPv6 traffic based on static entries in the IP Source Guard table, or dynamic entries in the DHCP Snooping table	379
Port Configuration	Enables IPv6 source guard and selects filter type per port	379
Static Binding		381
Add	Adds a static addresses to the source-guard binding table	381
Show	Shows static addresses in the source-guard binding table	381
Dynamic Binding	Displays the source-guard binding table for a selected interface	384
ARP Inspection		385
Configure General	Enables inspection globally, configures validation of additional address components, and sets the log rate for packet inspection	386
Configure VLAN	Enables ARP inspection on specified VLANs	388
Configure Interface	Sets the trust mode for ports, and sets the rate limit for packet inspection	389
Show Information		
Show Statistics	Displays statistics on the inspection process	391
Show Log	Shows the inspection log list	392
Application Filter	Discards CDP or PVST packets	393

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Administration		395
Log		396
System		396
Configure Global	Stores error messages in local memory	396
Show System Logs	Shows logged error messages	396
Remote	Configures the logging of messages to a remote logging process	398
SMTP	Sends an SMTP client message to a participating server	399
LLDP		401
Configure Global	Configures global LLDP timing parameters	401
Configure Interface		
Configure General	Sets the message transmission mode; enables SNMP notification; and sets the LLDP attributes to advertise	403
Add CA-Type	Specifies the physical location of the device attached to an interface	407
Modify CA-Type	Modifies the physical location of the device attached to an interface	407
Show CA-Type	Shows the physical location of the device attached to an interface	407
Show Local Device Information		409
General	Displays general information about the local device	409
Port/Trunk	Displays information about each interface	409
Port/Trunk Details	Displays detailed information about a local device connected to this switch	409
Show Remote Device Information		413
Port/Trunk	Displays information about a remote device connected to a port on this switch	413
Port/Trunk Details	Displays detailed information about a remote device connected to this switch	413
Show Device Statistics		421
General	Displays statistics for all connected remote devices	421
Port/Trunk	Displays statistics for remote devices on a selected port or trunk	421
PoE*	Power over Ethernet	423
PSE	Configures the maximum PoE power for the switch, and port power parameters for priority and maximum power allocation	424
SNMP	Simple Network Management Protocol	426
Configure Global	Enables SNMP agent status, and sets related trap functions	428
Configure Community		429
Add Community	Configures community strings and access mode	429
Show Community	Shows community strings and access mode	429

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Configure Engine		431
Set Engine ID	Sets the SNMP v3 engine ID on this switch	431
Add Remote Engine	Sets the SNMP v3 engine ID for a remote device	432
Show Remote Engine	Shows configured engine ID for remote devices	432
Configure View		433
Add View	Adds an SNMP v3 view of the OID MIB	433
Show View	Shows configured SNMP v3 views	433
Add OID Subtree	Specifies a part of the subtree for the selected view	433
Show OID Subtree	Shows the subtrees assigned to each view	433
Configure Group		436
Add	Adds a group with access policies for assigned users	436
Show	Shows configured groups and access policies	436
Configure User		
Add SNMPv3 Local User	Configures SNMPv3 users on this switch	441
Show SNMPv3 Local User	Shows SNMPv3 users configured on this switch	441
Change SNMPv3 Local User Group	Assign a local user to a new group	441
Add SNMPv3 Remote User	Configures SNMPv3 users from a remote device	443
Show SNMPv3 Remote User	Shows SNMPv3 users set from a remote device	441
Configure Trap		446
Add	Configures trap managers to receive messages on key events that occur this switch	446
Show	Shows configured trap managers	446
Configure Notify Filter		450
Add	Creates an SNMP notification log	450
Show	Shows the configured notification logs	450
Show Statistics	Shows the status of SNMP communications	452
RMON	Remote Monitoring	454
Configure Global		
Add		
Alarm	Sets threshold bounds for a monitored variable	454
Event	Creates a response event for an alarm	457
Show		
Alarm	Shows all configured alarms	454
Event	Shows all configured events	457

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Configure Interface		
Add		
History	Periodically samples statistics on a physical interface	459
Statistics	Enables collection of statistics on a physical interface	462
Show		
History	Shows sampling parameters for each entry in the history group	459
Statistics	Shows sampling parameters for each entry in the statistics group	462
Show Details		
History	Shows sampled data for each entry in the history group	459
Statistics	Shows sampled data for each entry in the history group	462
Cluster		464
Configure Global	Globally enables clustering for the switch; sets Commander status	465
Configure Member	Adds switch members to the cluster	466
Add	Adds candidate members to the cluster	466
Show	Shows the cluster members	466
Show Candidate	Shows candidate members	466
Show Member	Shows cluster switch member; managed switch members	468
Time Range	Configures the time to apply an ACL or PoE port	469
Add	Specifies the name of a time range	469
Show	Shows the name of configured time ranges	469
Add Rule		469
Absolute	Sets exact time or time range	469
Periodic	Sets a recurrent time	469
ERPS	Ethernet Ring Protection Switching	471
Configure Global	Activates ERPS globally	476
Configure Domain		476
Add	Creates an ERPS ring	476
Show	Shows list of configured ERPS rings, status, and settings	476
Configure Details	Configures ring parameters	476
Configure Operation	Blocks a ring port using Forced Switch or Manual Switch commands	492
CFM	Connectivity Fault Management	505
Configure Global	Configures global settings, including administrative status, cross-check start delay, link trace, and SNMP traps	509
Configure Interface	Configures administrative status on an interface	512

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Configure MD	Configure Maintenance Domains	512
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	512
Configure Details	Configures the archive hold time and fault notification settings	512
Show	Shows list of configured maintenance domains	512
Configure MA	Configure Maintenance Associations	517
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	517
Configure Details	Configures detailed settings, including continuity check status and interval level, cross-check status, and alarm indication signal parameters	517
Show	Shows list of configured maintenance associations	517
Configure MEP	Configures Maintenance End Points	521
Add	Configures MEPs at the domain boundary to provide management access for each maintenance association	521
Show	Shows list of configured maintenance end points	521
Configure Remote MEP	Configures Remote Maintenance End Points	523
Add	Configures a static list of remote MEPs for comparison against the MEPs learned through continuity check messages	523
Show	Shows list of configured remote maintenance end points	523
Transmit Link Trace	Sends link trace messages to isolate connectivity faults by tracing the path through a network to the designated target node	525
Transmit Loopback	Sends loopback messages to isolate connectivity faults by requesting a target node to echo the message back to the source	527
Transmit Delay Measure	Sends periodic delay-measure requests to a specified MEP within a maintenance association	529
Show Information		
Show Local MEP	Shows the MEPs configured on this device	531
Show Local MEP Details	Displays detailed CFM information about a specified local MEP in the continuity check database	532
Show Local MIP	Shows the MIPs on this device discovered by the CFM protocol	534
Show Remote MEP	Shows MEPs located on other devices which have been discovered through continuity check messages, or statically configured in the MEP database	535
Show Remote MEP Details	Displays detailed CFM information about a specified remote MEP in the continuity check database	536
Show Link Trace Cache	Shows information about link trace operations launched from this device	538
Show Fault Notification Generator	Displays configuration settings for the fault notification generator	539
Show Continuity Check Error	Displays CFM continuity check errors logged on this device	540

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
OAM		
Interface	Enables OAM on specified port, sets the mode to active or passive, and enables the reporting of critical events or errored frame events	496
Counters	Displays statistics on OAM PDUs	499
Event Log	Displays the log for recorded link events	500
Remote Interface	Displays information about attached OAM-enabled devices	501
Remote Loopback	Performs a loopback test on the specified port	502
UDLD	UniDirectional Link Detection	550
Configure Global	Configures the message probe interval, detection interval, and recovery interval	551
Configure Interface	Enables UDLD and aggressive mode which reduces the shut-down delay after loss of bidirectional connectivity is detected	552
Show Information	Displays UDLD neighbor information, including neighbor state, expiration time, and protocol intervals	554
LDB	Loopback Detection	555
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	556
Configure Interface	Enables loopback detection per interface	558
Tools		629
Ping	Sends ICMP echo request packets to another node on the network	629
Trace Route	Shows the route packets take to the specified destination	631
ARP	Address Resolution Protocol	632
Configure General	Sets the protocol timeout, and enables or disables proxy ARP for the specified VLAN	633
Configure Static Address		634
Add	Statically maps a physical address to an IP address	634
Show	Shows the MAC to IP address static table	634
Show Information		
ARP Address	Shows dynamically learned entries in the IP routing table, or internal addresses used by the switch	636
Statistics	Shows statistics for ARP messages crossing all interfaces on this router	636
IP		639
General		
Routing Interface		
Add Address	Configures an IP interface for a VLAN	639
Show Address	Shows the IP interfaces assigned to a VLAN	639

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Routing		
Static Routes		665
Add	Configures static routing entries	665
Show	Shows static routing entries	665
Routing Table		666
Show Information	Shows all routing entries, including local, static and dynamic routes	666
IPv6 Configuration		643
Configure Global	Sets an IPv6 default gateway for traffic with no known next hop	643
Configure Interface	Configures IPv6 interface address using auto-configuration or link-local address, and sets related protocol settings	644
Add IPv6 Address	Adds an global unicast, EUI-64, or link-local IPv6 address to an interface	649
Show IPv6 Address	Show the IPv6 addresses assigned to an interface	652
Show IPv6 Neighbor Cache	Displays information in the IPv6 neighbor discovery cache	653
Show Statistics		654
IPv6	Shows statistics about IPv6 traffic	654
ICMPv6	Shows statistics about ICMPv6 messages	654
UDP	Shows statistics about UDP messages	654
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	660
IP Service		669
DNS	Domain Name Service	669
General		669
Configure Global	Enables DNS lookup; defines the default domain name appended to incomplete host names	669
Add Domain Name	Defines a list of domain names that can be appended to incomplete host names	670
Show Domain Names	Shows the configured domain name list	670
Add Name Server	Specifies IP address of name servers for dynamic lookup	672
Show Name Servers	Shows the name server address list	672
Static Host Table		673
Add	Configures static entries for domain name to address mapping	673
Show	Shows the list of static mapping entries	673
Modify	Modifies the static address mapped to the selected host name	673
Cache	Displays cache entries discovered by designated name servers	674

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
DHCP	Dynamic Host Configuration Protocol	675
Client	Specifies the DHCP client identifier for an interface	675
Relay	Specifies DHCP relay servers	677
Dynamic Provision	Enables dynamic provisioning via DHCP	678
PPPoE Intermediate Agent		679
Configure Global	Enables PPPoE IA on the switch, sets access node identifier, sets generic error message	679
Configure Interface	Enables PPPoE IA on an interface, sets trust status, enables vendor tag stripping, sets circuit ID and remote ID	681
Show Statistics	Shows statistics on PPPoE IA protocol messages	682
Multicast		559
IGMP Snooping		560
General	Enables multicast filtering; configures parameters for IPv4 multicast snooping	562
Multicast Router		566
Add Static Multicast Router	Assigns ports that are attached to a neighboring multicast router	566
Show Static Multicast Router	Displays ports statically configured as attached to a neighboring multicast router	566
Show Current Multicast Router	Displays ports attached to a neighboring multicast router, either through static or dynamic configuration	566
IGMP Member		568
Add Static Member	Statically assigns multicast addresses to the selected VLAN	568
Show Static Member	Shows multicast addresses statically configured on the selected VLAN	568
Interface		570
Configure VLAN	Configures IGMP snooping per VLAN interface	570
Show VLAN Information	Shows IGMP snooping settings per VLAN interface	570
Configure Interface	Configures the interface to drop IGMP query packets or all multicast data packets	576
Forwarding Entry	Displays the current multicast groups learned through IGMP Snooping	577
Filter		582
Configure General	Enables IGMP filtering for the switch	583
Configure Profile		583
Add	Adds IGMP filter profile; and sets access mode	583
Show	Shows configured IGMP filter profiles	583
Add Multicast Group Range	Assigns multicast groups to selected profile	583
Show Multicast Group Range	Shows multicast groups assigned to a profile	583
Configure Interface	Assigns IGMP filter profiles to port interfaces and sets throttling action	586

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Statistics		578
Show Query Statistics	Shows statistics for query-related messages	578
Show VLAN Statistics	Shows statistics for protocol messages and number of active groups	578
Show Port Statistics	Shows statistics for protocol messages and number of active groups	578
Show Trunk Statistics	Shows statistics for protocol messages and number of active groups	578
MLD Snooping		587
General	Enables multicast filtering; configures parameters for IPv6 multicast snooping	588
Interface	Configures Immediate Leave status for a VLAN	589
Multicast Router		590
Add Static Multicast Router	Assigns ports that are attached to a neighboring multicast router	590
Show Static Multicast Router	Displays ports statically configured as attached to a neighboring multicast router	590
Show Current Multicast Router	Displays ports attached to a neighboring multicast router, either through static or dynamic configuration	590
MLD Member		592
Add Static Member	Statically assigns multicast addresses to the selected VLAN	592
Show Static Member	Shows multicast addresses statically configured on the selected VLAN	592
Show Current Member	Shows multicast addresses associated with the selected VLAN, either through static or dynamic configuration	592
Group Information	Displays known multicast groups, member ports, the means by which each group was learned, and the corresponding source list	594
MVR	Multicast VLAN Registration	596
Configure Global	Configures proxy switching and robustness value	597
Configure Domain	Enables MVR for a domain, sets the MVR VLAN, forwarding priority, and upstream source IP	599
Configure Profile		600
Add	Configures multicast stream addresses	600
Show	Shows multicast stream addresses	600
Associate Profile		600
Add	Maps an address profile to a domain	600
Show	Shows addresses profile to domain mapping	600
Configure Interface	Configures MVR interface type and immediate leave mode; also displays MVR operational and active status	603
Configure Static Group Member		605
Add	Statically assigns MVR multicast streams to an interface	605
Show	Shows MVR multicast streams assigned to an interface	605

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show Member	Shows the multicast groups assigned to an MVR VLAN, the source address of the multicast services, and the interfaces with active subscribers	607
Show Statistics		608
Show Query Statistics	Shows statistics for query-related messages	608
Show VLAN Statistics	Shows statistics for protocol messages and number of active groups	608
Show Port Statistics	Shows statistics for protocol messages and number of active groups	608
Show Trunk Statistics	Shows statistics for protocol messages and number of active groups	608
MVR6	Multicast VLAN Registration for IPv6	612
Configure Global	Configures proxy switching and robustness value	613
Configure Domain	Enables MVR for a domain, sets the MVR VLAN, forwarding priority, and upstream source IP	615
Configure Profile		616
Add	Configures multicast stream addresses	616
Show	Shows multicast stream addresses	616
Associate Profile		616
Add	Maps an address profile to a domain	616
Show	Shows addresses profile to domain mapping	616
Configure Interface	Configures MVR interface type and immediate leave mode; also displays MVR operational and active status	619
Configure Port	Configures MVR attributes for a port	619
Configure Trunk	Configures MVR attributes for a trunk	619
Configure Static Group Member		621
Add	Statically assigns MVR multicast streams to an interface	621
Show	Shows MVR multicast streams assigned to an interface	621
Show Member	Shows the multicast groups assigned to an MVR VLAN, the source address of the multicast services, and the interfaces with active subscribers	623
Show Statistics		624
Show Query Statistics	Shows statistics for query-related messages	624
Show VLAN Statistics	Shows statistics for protocol messages and number of active groups	624
Show Port Statistics	Shows statistics for protocol messages and number of active groups	624
Show Trunk Statistics	Shows statistics for protocol messages and number of active groups	624

* ECS4120-28P

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 4: System Information

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

System Description	ECS4120-28T
System Object ID	1.3.6.1.4.1.259.10.1.45.101
System Up Time	0 days, 2 hours, 22 minutes, and 57.62 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.
- ◆ **Redundant Power Status** – Displays the status of the redundant power supply.

Management Software Information

- ◆ **Role** – Shows that this switch is operating as Master or Slave.
- ◆ **EPLD Version** – Version number of EEPROM Programmable Logic Device.
- ◆ **Loader Version** – Version number of loader code.
- ◆ **Operation Code Version** – Version number of runtime code.

Thermal Sensors


- ◆ **Thermal Detector** – The first detector is near the air flow intake vents. The second detector is near the switch ASIC and CPU.
- ◆ **Temperature** – Temperature at specified thermal detection point.

Web Interface

To view hardware and software version information.

- 1. Click System, then Switch.

Figure 5: General Switch Information



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 6: 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 241.](#))
- ◆ **Static Entry Individual Port** – This switch allows static filtering for unicast and multicast addresses. (Refer to [“Setting Static Addresses” on page 201.](#))
- ◆ **VLAN Version Number** – Based on IEEE 802.1Q, “1” indicates Bridges that support only single spanning tree (SST) operation, and “2” indicates Bridges that support multiple spanning tree (MST) operation.
- ◆ **VLAN Learning** – This switch uses Independent VLAN Learning (IVL), where each port maintains its own filtering database.
- ◆ **Local VLAN Capable** – This switch does not support multiple local bridges outside of the scope of 802.1Q defined VLANs.

- ◆ **Configurable PVID Tagging** – This switch allows you to override the default Port VLAN ID (PVID used in frame tags) and egress status (VLAN-Tagged or Untagged) on each port. (Refer to [“VLAN Configuration” on page 159.](#))
- ◆ **Max Supported VLAN Numbers** – The maximum number of VLANs supported on this switch.
- ◆ **Max Supported VLAN ID** – The maximum configurable VLAN identifier supported on this switch.

Web Interface

To view Bridge Extension information:

1. Click System, then Capability.

Figure 7: Displaying Bridge Extension Configuration

The screenshot displays the 'System > Capability' configuration page. It features a 'General Capability' section with a 'Jumbo Frame' toggle set to 'Enabled'. Below this is the 'Bridge Extension' section, which contains a list of configuration items and their current values. At the bottom right of the configuration area are 'Apply' and 'Revert' buttons.

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

Apply Revert

Managing System Files

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

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

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

Command Usage

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

Parameters

The following parameters are displayed:

- ◆ **Copy Type** – The firmware copy operation includes these options:
 - HTTP Upload – Copies a file from a management station to the switch.
 - HTTP Download – Copies a file from the switch to a management station
 - TFTP Upload – Copies a file from a TFTP server to the switch.
 - TFTP Download – Copies a file from the switch to a TFTP server.
 - FTP Upload – Copies a file from an FTP server to the switch.
 - FTP Download – Copies a file from the switch to an FTP server.
- ◆ **FTP/TFTP Server IP Address** – The IP address of an FTP/TFTP server.
- ◆ **User Name** – The user name for FTP server access.
- ◆ **Password** – The password for FTP server access.
- ◆ **File Type** – Specify Operation Code to copy firmware or Config File to copy configuration settings.
- ◆ **File Name** – The file name should not contain slashes (\ or /), the leading letter of the file name should not be a period (.), and the maximum length for file

names is 32 characters for files on the switch or 127 characters for files on the server. (Valid characters: A-Z, a-z, 0-9, "-", "_")



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

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

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

Web Interface

To copy firmware files:

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

Figure 8: Copy Firmware

The screenshot shows the 'System > File' web interface. At the top, there is a breadcrumb 'System > File'. Below it, the 'Action' dropdown is set to 'Copy'. The 'Copy Type' dropdown is set to 'TFTP Upload'. The 'TFTP Server IP Address' field contains '192.168.2.99'. The 'File Type' dropdown is set to 'Operation Code'. The 'Source File Name' field contains 'ecs4120-run_v1.0.2.3.bix'. The 'Destination File Name' section has two radio buttons: the first is for 'ecs4120-run_v115.12.31.1.bix' and the second, which is selected, is for 'ecs4120-run_v1.0.2.3.bix'. 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 9: Saving the Running Configuration

System > File

Action: 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 10: Setting Start-Up Files

System > File

Action: Set Start-Up

File List Total: 4

	File Name	File Type	Start-Up	Modify Time	Size (bytes)
☐	ecs4120-run_v1.0.2.0.bix	Operation Code	N	2015-12-31 04:56:50	18815572
☒	ecs4120-run_v115.12.31.1.bix	Operation Code	Y	2015-08-07 07:00:16	18958932
☐	Factory_Default_Config.cfg	Config File	N	2015-08-07 03:59:12	455
☒	startup1.cfg	Config File	Y	2015-08-07 03:59:24	1537

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 11: Displaying System Files

File Name	File Type	Start-Up	Modify Time	Size (bytes)
ecs4120-run_v1.0.2.0.bix	Operation Code	N	2015-12-31 04:56:50	18815572
ecs4120-run_v115.12.31.1.bix	Operation Code	Y	2015-08-07 07:00:16	18958932
Factory_Default_Config.cfg	Config File	N	2015-08-07 03:59:12	455
startup1.cfg	Config File	Y	2015-08-07 03:59:24	1537

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 *ECS4120-series.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 *ECS4120-Series.BIX* from the server even though *ECS4120-series.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, *ecs4120-series.bix* and *ECS4120-Series.bix* are considered to be unique files. Thus, if the upgrade file is stored as *ECS4120-Series.bix* (or even *EcS4120-Series.bix*) on a case-sensitive server, then the switch (requesting *ecs4120-series.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 *ECS4120-series.bix* filename must not be included since it is automatically appended by the switch. (Options: ftp, tftp)

The following syntax must be observed:

tftp://host[/filedir]/

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

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

- **ftp://** – Defines FTP protocol for the server connection.
- *username* – Defines the user name for the FTP connection. If the user name is omitted, then "anonymous" is the assumed user name for the connection.

If no user name nor password is required for the connection, then the "@" character cannot be used in the path name.
- *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 12: Configuring Automatic Code Upgrade

System > File

Action: Automatic Operation Code Upgrade ▼

Automatic Opcode Upgrade ☒ Enabled

Automatic Upgrade Location URL

Note: For automatic upgrades, the operation code file name must be set as ECS4120-Series.bix.

Apply 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 1.2.1.3; new version 1.2.1.6
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 13: Manually Setting the System Clock

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

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

Parameters

The following parameters are displayed:

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

Web Interface

To set the polling interval for SNTP:

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

Figure 14: Setting the Polling Interval for SNTP

The screenshot shows the 'System > Time' configuration page. At the top, there is a 'Step:' dropdown menu set to '1. Configure General'. Below this, the 'Current Time' is displayed as '2014-5-30 9:59:20'. The 'Maintain Type' is set to 'SNTP' via a dropdown menu. Under the 'SNTP Configuration' section, the 'SNTP Polling Interval (16-16384)' is set to '16' seconds. At the bottom right, there are 'Apply' and 'Revert' buttons.

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

Parameters

The following parameters are displayed:

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

You can enable NTP authentication to ensure that reliable updates are received from only authorized NTP servers. The authentication keys and their associated key number must be centrally managed and manually distributed to NTP servers and clients. The key numbers and key values must match on both the server and client.

- ◆ **Polling Interval** – Shows the interval between sending requests for a time update from NTP servers. (Fixed: 1024 seconds)

Web Interface

To set the clock maintenance type to NTP:

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

Figure 15: Configuring NTP

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

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

Specifying SNTP Time Servers

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

Parameters

The following parameters are displayed:

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

Web Interface

To set the SNTP time servers:

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

Figure 16: Specifying SNTP Time Servers

The screenshot shows a web interface for configuring system time. The breadcrumb is 'System > Time'. Below it, there are two dropdown menus: 'Step' set to '2. Configure Time Server' and 'Action' set to 'Configure SNTP Server'. Below these are three text input fields labeled 'SNTP Server IP Address 1', 'SNTP Server IP Address 2', and 'SNTP Server IP Address 3'. The first field contains '10.1.0.19', the second contains '137.62.140.80', and the third contains '128.250.36.2'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

Specifying NTP Time Servers

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

Parameters

The following parameters are displayed:

- ◆ **NTP Server IP Address** – Adds the IPv4 or IPv6 address for up to 50 time servers. The switch will poll the specified time servers for updates when the clock maintenance type is set to NTP on the System > Time (Configure General) page. It issues time synchronization requests at a fixed interval of 1024 seconds. The switch will poll all the time servers configured, the responses received are filtered and compared to determine the most reliable and accurate time update for the switch.
- ◆ **Version** – Specifies the NTP version supported by the server. (Fixed: Version 3)

- ◆ **Authentication Key** – Specifies the number of the key in the NTP Authentication Key List to use for authentication with the configured server. NTP authentication is optional. If enabled on the System > Time (Configure General) page, you must also configure at least one key on the System > Time (Add NTP Authentication Key) page. (Range: 1-65535)

Web Interface

To add an NTP time server to the server list:

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

Figure 17: Adding an NTP Time Server

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

To show the list of configured NTP time servers:

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

Figure 18: Showing the NTP Time Server List

The screenshot shows the 'System > Time' configuration page with the 'Action:' dropdown set to 'Show NTP Server'. Below the form, there's a table titled 'NTP Server List' with a 'Total: 1' indicator. The table has four columns: a checkbox, 'Server IP Address', 'Version', and 'Authentication Key'. It contains one row with the IP '192.168.3.20', version '3', and key '3'. At the bottom right, there are 'Delete' and 'Revert' buttons.

	Server IP Address	Version	Authentication Key
☐	192.168.3.20	3	3

Specifying NTP Authentication Keys

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

Parameters

The following parameters are displayed:

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

NTP authentication key numbers and values must match on both the server and client.

Web Interface

To add an entry to NTP authentication key list:

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


Figure 19: Adding an NTP Authentication Key

System > Time	
Step:	2. Configure Time Server
Action:	Add NTP Authentication Key
Authentication Key (1-65535)	3
Key Context (1-32)	S1507N122103J068173M
Apply Revert	

To show the list of configured NTP authentication keys:

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

Figure 20: 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 21: Setting the Time Zone

System > Time

Step: 3. Configure Time Zone ▼

☐ Predefined Configuration

☒ User Defined Configuration

Direction: After UTC ▼

Name: UTC

Hours (0-13): 0

Minutes (0-59): 0

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

Apply Revert

Configuring Summer Time

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

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

Parameters

The following parameters are displayed in the web interface:

General Configuration

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

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

Table 5: Predefined Summer-Time Parameters

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

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

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

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

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

Web Interface

To specify summer time settings:

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

Figure 22: Configuring Summer Time

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

- Summer Time in Effect:** A radio button labeled 'No' is selected.
- Status:** An 'Enabled' checkbox is present and unchecked.
- Name:** A text input field is empty.
- Mode:** A dropdown menu is set to 'Predefined'.
- Predefined Mode Configuration:** A section header.
- Daylight Savings:** A dropdown menu is set to 'Australia'.

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

Configuring the Console Port

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

Parameters

The following parameters are displayed:

- ◆ **Login Timeout** – Sets the interval that the system waits for a user to log into the CLI. If a login attempt is not detected within the timeout interval, the connection is terminated for the session. (Range: 10-300 seconds; Default: 300 seconds)
- ◆ **Exec Timeout** – Sets the interval that the system waits until user input is detected. If user input is not detected within the timeout interval, the current session is terminated. (Range: 60-65535 seconds; Default: 600 seconds)
- ◆ **Password Threshold** – Sets the password intrusion threshold, which limits the number of failed logon attempts. When the logon attempt threshold is reached, the system interface becomes silent for a specified amount of time (set by the Silent Time parameter) before allowing the next logon attempt. (Range: 1-120; Default: 3 attempts)
- ◆ **Silent Time** – Sets the amount of time the management console is inaccessible after the number of unsuccessful logon attempts has been exceeded. (Range: 1-65535 seconds; Default: Disabled)
- ◆ **Data Bits** – Sets the number of data bits per character that are interpreted and generated by the console port. If parity is being generated, specify 7 data bits

per character. If no parity is required, specify 8 data bits per character.
(Default: 8 bits)

- ◆ **Stop Bits** – Sets the number of the stop bits transmitted per byte.
(Range: 1-2; Default: 1 stop bit)
- ◆ **Parity** – Defines the generation of a parity bit. Communication protocols provided by some terminals can require a specific parity bit setting. Specify Even, Odd, or None. (Default: None)
- ◆ **Speed** – Sets the terminal line's baud rate for transmit (to terminal) and receive (from terminal). Set the speed to match the baud rate of the device connected to the serial port. (Range: 9600, 19200, 38400, 57600, or 115200 baud; Default: 115200 baud)



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

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

Web Interface

To configure parameters for the console port:

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

Figure 23: Console Port Settings

System > Console

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

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 24: Telnet Connection Settings

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

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

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

Displaying CPU Utilization

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

Parameters

The following parameters are displayed:


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

Web Interface

To display CPU utilization:

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

Figure 25: 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: 50-500 pps; Default: 500 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: 50-500 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 26: Configuring CPU Guard

The screenshot shows the 'System > CPU Guard' configuration page. It contains several settings: 'CPU Guard Status' with an 'Enabled' checkbox; 'High Watermark (40-100)' set to 90%; 'Low Watermark (40-100)' set to 70%; 'Maximum Threshold (50-500)' set to 500 packets/sec; 'Minimum Threshold (50-500)' set to 50 packets/sec; 'Trap Status' with an 'Enabled' checkbox; and 'Current Threshold' set to 500 packets/sec. At the bottom right are 'Apply' and 'Revert' buttons.

System > CPU Guard	
CPU Guard Status	☐ Enabled
High Watermark (40-100)	90 %
Low Watermark (40-100)	70 %
Maximum Threshold (50-500)	500 packets/sec
Minimum Threshold (50-500)	50 packets/sec
Trap Status	☐ Enabled
Current Threshold	500 packets/sec
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 27: Displaying Memory Utilization

System > Memory Status		
Memory Status		
Free Size	265,756,672 bytes	49%
Used Size	271,114,240 bytes	51%
Total	536,870,912 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 79](#)).

Parameters

The following parameters are displayed:

System Reload Information

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

System Reload Configuration

- ◆ **Reset Mode** – Restarts the switch immediately or at the specified time(s).

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

Time

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

Period

- Daily - Every day.
- Weekly - Day of the week at which to reload.
(Range: Sunday ... Saturday)
- Monthly - Day of the month at which to reload. (Range: 1-31)

Web Interface

To restart the switch:

1. Click System, then Reset.
2. Select the required reset mode.
3. For any option other than to reset immediately, fill in the required parameters
4. Click Apply.

5. When prompted, confirm that you want reset the switch.

Figure 28: Restarting the Switch (Immediately)

System > Reset

System Reload Information:
No configured settings for reloading.

System Reload Configuration:
Reset Mode Immediately

Message from webpage
Do you want to reset the switch immediately?

Figure 29: Restarting the Switch (In)

System > Reset

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

System Reload Configuration:
Reset Mode In

Reload switch in 1 hours 30 minutes.

Note: The specified time must be equal to or less than 24 days.

Figure 30: Restarting the Switch (At)

System > Reset

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

RefreshCancel

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

ApplyRevert

Figure 31: Restarting the Switch (Regularly)

System > Reset

System Reload Information:
No configured settings for reloading.

RefreshCancel

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

ApplyRevert

Interface Configuration

This chapter describes the following topics:

- ◆ [Port Configuration](#) – Configures connection settings, including auto-negotiation, or manual setting of speed, duplex mode, and flow control.
- ◆ [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.
- ◆ [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 1000BASE-T standard does not support forced mode. Auto-negotiation should always be used to establish a connection over any 1000BASE-T port or trunk. If not used, the success of the link process cannot be guaranteed when connecting to other types of switches.
- ◆ When using auto-negotiation¹, the optimal settings will be negotiated between the link partners based on their advertised capabilities. To set flow control and symmetric pause frames under auto-negotiation, the required operation modes must be specified in the capabilities list for an interface.

When auto-negotiation is enabled¹, the only attributes which can be advertised include speed, duplex mode, flow control and symmetric pause frames.
- ◆ Although not covered in the 1000BASE-T standard, auto-negotiation on 1000BASE-T ports must be disabled before you can force an RJ-45 interface to use the Speed/Duplex mode or Flow Control options. Note that Speed/Duplex mode is fixed at 100full for 100BASE-FX transceivers, 1000full for Gigabit transceivers, and 10Gfull for 10 Gigabit transceivers.

Using Jumbo Frames

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

1. Support for auto-negotiation depends on port or transceiver type, such as 1000BASE-T or 1G SFP.

- ◆ For other traffic types, calculation of overall frame size is basically the same, including the additional header fields SA(6) + DA(6) + Type(2) + VLAN-Tag(4) for tagged packets, (for untagged packets, the 4-byte field will not be added by switch), and the payload. This should all be less than the configured port MTU, including the CRC at the end of the frame.
- ◆ For QinQ, the overall frame size is still calculated as described above, but does not add the length of the second tag to the frame.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28/52)
- ◆ **Type** – Indicates the port type. (1000BASE-T, 100BASE-FX SFP, 1000BASE SFP, 10GBASE 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 forced transceiver mode for SFP/SFP+ ports, or forced/preferred port type for RJ-45/SFP combination ports.
 - **None** - Forced transceiver mode is not used for SFP/SFP+ ports. (This is the default setting for RJ-45 ports and SFP/SFP+ ports.)
 - **Copper-Forced** - Always uses the RJ-45 port. (Only applies to combination RJ-45/SFP ports 21-24 on the ECS4120-28F.)
 - **SFP-Forced 1000SFP** - Always uses the SFP/SFP+ port at 1000 Mbps, Full Duplex.
 - **SFP-Forced 100FX** - Always uses the SFP port at 100 Mbps, Full Duplex. (Only applies to SFP ports on the ECS4120-28F/28F-DC.)
 - **SFP-Forced 10GSFP** - Always uses the SFP+ port at 10 Gbps, Full Duplex.
- ◆ **Autonegotiation** (Port Capabilities) – Allows auto-negotiation to be enabled/disabled. When auto-negotiation is enabled¹, you need to specify the capabilities to be advertised. When auto-negotiation is disabled, you can force the setting for speed, duplex mode, and flow control. The following capabilities are supported.
 - **10h** - Supports 10 Mbps half-duplex operation.
 - **10f** - Supports 10 Mbps full-duplex operation.

- **100h** - Supports 100 Mbps half-duplex operation.
- **100f** - Supports 100 Mbps full-duplex operation.
- **1000f** - Supports 1000 Mbps full-duplex operation.
- **10Gf** (10G SFP+ ports only) - Supports 10 Gbps full-duplex operation.
- **Sym** (Symmetric) - Check this item to transmit and receive pause frames.
- **FC** (Flow Control) - Flow control can eliminate frame loss by “blocking” traffic from end stations or segments connected directly to the switch when its buffers fill. When enabled, back pressure is used for half-duplex operation and IEEE 802.3-2005 (formally IEEE 802.3x) for full-duplex operation.

Default: Autonegotiation enabled on 1000BASE-T and 1000BASE-SFP, disabled on 10GBASE;

Advertised capabilities² for -

100BASE-FX (SFP) – 100full

1000BASE-T – 10half, 10full, 100half, 100full, 1000full

1000BASE-SX/LX/LHX/ZX (SFP³ / SFP+) – 1000full

10GBASE- CR/SR/LR/ER (SFP+) – 10Gfull

- ◆ **Speed/Duplex** – Allows you to manually set the port speed and duplex mode. (i.e., with auto-negotiation disabled)
- ◆ **Flow Control** – Allows automatic or manual selection of flow control. (Default: Enabled)
- ◆ **MTU Size** – The maximum transfer unit (MTU) allowed for layer 2 packets crossing a 1G/10G Ethernet port or trunk. (Range: 1500-9216 bytes; Default: 1518 bytes)
- ◆ **Link Up Down Trap** – Issues a notification message whenever a port link is established or broken. (Default: Enabled)

Web Interface

To configure port connection parameters:

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

2. This list only includes tested Edgecore transceiver types.

3. These items are only applicable to the ECS4120-28F.

Figure 32: Configuring Connections by Port List

Interface > Port > General

Action: Configure by Port List

Port List Total: 28

Port	Type	Name	Admin	Media Type	Autonegotiation	Speed Duplex	Flow Control	MTU Size (1500-9216)	Link Up Down Trap
1	1000BASE-T		☒ Enabled	None	☒ Enabled ☒ 10h ☒ 100h ☒ 1000f ☐ FC ☒ 10f ☒ 100f ☐ Sym	100full	☐ Enabled	1518	☒ Enabled
2	1000BASE-T		☒ Enabled	None	☒ Enabled ☒ 10h ☒ 100h ☒ 1000f ☐ FC ☒ 10f ☒ 100f ☐ Sym	100full	☐ Enabled	1518	☒ Enabled
3	1000BASE-T		☒ Enabled	None	☒ Enabled ☒ 10h ☒ 100h ☒ 1000f ☐ FC ☒ 10f ☒ 100f ☐ Sym	100full	☐ Enabled	1518	☒ Enabled

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

For more information on command usage and a description of the parameters, refer to [“Configuring by Port List” on page 106](#).

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 33: Configuring Connections by Port Range

Interface > Port > General

Action: Configure by Port Range ▼

Port Range (1-28)

Admin ☒ Enabled

Autonegotiation ☒ Enabled

☐ 10h ☐ 100h ☐ 1000f ☐ Sym

☐ 10f ☐ 100f ☐ 10Gf ☐ FC

Speed Duplex 10half ▼

Flow Control ☒ Enabled

MTU Size (1500-9216)

Link Up Down Trap ☒ 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, 1000BASE SFP, 10GBASE 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+ ports, or forced port type for RJ-45/SFP combination ports used in the ECS4120-28F. (None, Copper-Forced⁴, SFP-Forced 1000SFP, SFP-Forced 100FX⁴, SFP-Forced 10GSFP)
- ◆ **Autonegotiation** – Shows if auto-negotiation is enabled or disabled.
- ◆ **Oper Speed Duplex** – Shows the current speed and duplex mode.

4. These items only apply to the ECS4120-28F.


- ◆ **Oper Flow Control** – Shows the flow control type used.
- ◆ **MTU Size** – The maximum transfer unit (MTU) allowed for layer 2 packets crossing a Gigabit or 10 Gigabit Ethernet port or trunk.
- ◆ **Link Up Down Trap** – Shows if a notification message will be sent whenever a port link is established or broken. (Default: Enabled)

Web Interface

To display port connection parameters:

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

Figure 34: Displaying Port Information



Port	Type	Name	Admin	Oper Status	Shutdown Reason	Media Type	Autonegotiation	Oper Speed Duplex	Oper Flow Control	MTU Size	Link Up Down Trap
1	1000BASE-T		Enabled	Up		None	Enabled	100full	None	1518	Enabled
2	1000BASE-T		Enabled	Down		None	Enabled	1000full	None	1518	Enabled
3	1000BASE-T		Enabled	Down		None	Enabled	1000full	None	1518	Enabled
4	1000BASE-T		Enabled	Down		None	Enabled	1000full	None	1518	Enabled
5	1000BASE-T		Enabled	Down		None	Enabled	1000full	None	1518	Enabled

Showing Port or Trunk Statistics

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



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

Parameters

These parameters are displayed:

Table 6: Port Statistics

Parameter	Description
<i>Interface Statistics</i>	
Received Octets	The total number of octets received on the interface, including framing characters.
Transmitted Octets	The total number of octets transmitted out of the interface, including framing characters.
Received Errors	The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.
Transmitted Errors	The number of outbound packets that could not be transmitted because of errors.
Received Unicast Packets	The number of subnetwork-unicast packets delivered to a higher-layer protocol.
Transmitted Unicast Packets	The total number of packets that higher-level protocols requested be transmitted to a subnetwork-unicast address, including those that were discarded or not sent.
Received Discarded Packets	The number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. One possible reason for discarding such a packet could be to free up buffer space.
Transmitted Discarded Packets	The number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could be to free up buffer space.
Received Multicast Packets	The number of packets, delivered by this sub-layer to a higher (sub-)layer, which were addressed to a multicast address at this sub-layer.
Transmitted Multicast Packets	The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a multicast address at this sub-layer, including those that were discarded or not sent.
Received Broadcast Packets	The number of packets, delivered by this sub-layer to a higher (sub-)layer, which were addressed to a broadcast address at this sub-layer.
Transmitted Broadcast Packets	The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a broadcast address at this sub-layer, including those that were discarded or not sent.
Received Unknown Packets	The number of packets received via the interface which were discarded because of an unknown or unsupported protocol.
<i>Etherlike Statistics</i>	
Single Collision Frames	The number of successfully transmitted frames for which transmission is inhibited by exactly one collision.
Multiple Collision Frames	A count of successfully transmitted frames for which transmission is inhibited by more than one collision.
Late Collisions	The number of times that a collision is detected later than 512 bit-times into the transmission of a packet.
Excessive Collisions	A count of frames for which transmission on a particular interface fails due to excessive collisions. This counter does not increment when the interface is operating in full-duplex mode.

Table 6: Port Statistics (Continued)

Parameter	Description
Deferred Transmissions	A count of frames for which the first transmission attempt on a particular interface is delayed because the medium was busy.
Frames Too Long	A count of frames received on a particular interface that exceed the maximum permitted frame size.
Symbol Errors	A count of the times there was an invalid data symbol when a valid carrier was present.
Alignment Errors	The number of alignment errors (missynchronized data packets).
FCS Errors	A count of frames received on a particular interface that are an integral number of octets in length but do not pass the FCS check. This count does not include frames received with frame-too-long or frame-too-short error.
SQE Test Errors	A count of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular interface.
Carrier Sense Errors	The number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame.
Internal MAC Receive Errors	A count of frames for which reception on a particular interface fails due to an internal MAC sublayer receive error.
Internal MAC Transmit Errors	A count of frames for which transmission on a particular interface fails due to an internal MAC sublayer transmit error.
Pause Frames Output	A count of MAC control frames transmitted on this interface with an opcode indicating the pause operation.
<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.

Table 6: Port Statistics (Continued)


Parameter	Description
64 Bytes Packets	The total number of packets (including bad packets) received and transmitted that were 64 octets in length (excluding framing bits but including FCS octets).
65-127 Byte Packets	The total number of packets (including bad packets) received and transmitted where the number of octets fall within the specified range (excluding framing bits but including FCS octets).
128-255 Byte Packets	
256-511 Byte Packets	
512-1023 Byte Packets	
1024-1518 Byte Packets	
1519-1536 Byte Packets	
<i>Utilization Statistics</i>	
Input Octets in kbits per second	Number of octets entering this interface in kbits/second.
Input Packets per second	Number of packets entering this interface per second.
Input Utilization	The input utilization rate for this interface.
Output Octets in kbits per second	Number of octets leaving this interface in kbits/second.
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 if you need to update the screen.


Figure 35: 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 36: 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](#).

Parameters

These parameters are displayed:

Add

- ◆ **Port** – Port number. (Range: 1-28/52)
- ◆ **History Name** – Name of sample interval. (Range: 1-32 characters)
- ◆ **Interval** – The interval for sampling statistics. (Range: 1-86400 seconds)
- ◆ **Requested Buckets** – The number of samples to take. (Range: 1-96)

Show

- ◆ **Port** – Port number. (Range: 1-28/52)
- ◆ **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: 1-28/52)
- ◆ **Name** – Name of sample interval.
- ◆ **Interval** – The interval for sampling statistics.

- ◆ **Requested Buckets** - The number of samples to take.
- ◆ **Granted Buckets** - The number of buckets granted.
- ◆ **Status** - Shows if this entry is active.

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 37: Configuring a History Sample

Interface > Port > History

Action: Add ▼

Port 1 ▼

History Name rd#1

Interval (1-86400) 60 seconds

Requested Buckets (1-96) 50

Apply Revert

To show the configured entries for a history sample:

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

Figure 38: Showing Entries for History Sampling

The screenshot shows the 'Interface > Port > History' page. At the top, there is a breadcrumb trail and a help icon. Below this, there is an 'Action:' dropdown menu set to 'Show'. Underneath, there is a 'Port' dropdown menu set to '1'. Below that, there is a 'History Name List' with a 'Total: 3' indicator. The main content is a table with the following data:

	History Name	Interval	Requested Buckets
☐	15min	900	96
☐	1day	86400	7
☐	rd#1	60	50

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

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 39: Showing Status of Statistical History Sample

The screenshot shows the 'Interface > Port > History' page. At the top, there is a breadcrumb trail and a help icon. Below this, there is an 'Action:' dropdown menu set to 'Show Details'. Underneath, there is a 'Mode' section with four radio buttons: 'Status' (selected), 'Current Entry', 'Input Previous Entries', and 'Output Previous Entries'. Below the 'Mode' section, there is a 'Port' dropdown menu set to '1' and a 'Name' dropdown menu set to '15min'. Below that, there is a 'History Status' section with the following data:

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

At the bottom right of the page, there is a 'Refresh' button.

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 40: Showing Current Statistics for a History Sample

Interface > Port > History

Action: Show Details

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

Port
1

Name
15min

Current Entry

Start Time	00d 08:45:09		
Received Octets	114377	Transmitted Octets	408498
Received Errors	0	Transmitted Errors	0
Received Unicast Packets	504	Transmitted Unicast Packets	589
Received Discarded Packets	0	Transmitted Discarded Packets	0
Received Multicast Packets	412	Transmitted Multicast Packets	25
Received Broadcast Packets	8	Transmitted Broadcast Packets	0
Received Unknown Packets	0		

Refresh

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 41: Showing Ingress Statistics for a History Sample

Interface > Port > History

Action: Show Details

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

Port: 1

Name: 15min

Input Previous Entry List Total: 16

Start Time	%	Octets	Unicast	Multicast	Broadcast	Discarded	Errors	Unknown Proto
00d 00:00:0	0.00	9401	0	40	36	24	0	0
00d 00:15:0	0.00	3246	0	30	9	9	0	0
00d 00:30:0	0.00	2999	0	30	8	8	0	0
00d 00:45:0	0.00	3863	0	30	17	17	0	0
00d 01:00:0	0.00	3511	0	29	14	14	0	0
00d 01:15:0	0.00	3246	0	30	9	9	0	0
00d 01:30:0	0.00	2999	0	30	8	8	0	0
00d 01:45:0	0.00	2999	0	30	8	8	0	0
00d 02:00:0	0.00	3575	0	30	14	14	0	0
00d 02:15:0	0.00	3246	0	30	9	9	0	0

Refresh

Transceiver Data and Thresholds

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: 25-28⁵/49-52)
- ◆ **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.

5. Due a chip limitation, transceiver data cannot be displayed on ports 1-24 on the ECS4120-28F.

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 42: Displaying Transceiver Data

Interface > Port > Transceiver

Port

11

General

Connector Type	LC
Fiber Type	Multimode 50um (M5), Multimode 62.5um (M6)
Eth Compliance Codes	1000BASE-SX
Baud Rate	2100 MBd
Vendor OUI	00-90-85
Vendor Name	FINISAR CORP.
Vendor PN	FTLF8519P3BTL
Vendor Rev	A
Vendor SN	PKM1XUU
Date Code	11-05-25

DDM Information

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

Configuring Transceiver Thresholds

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

Parameters

These parameters are displayed:

- ◆ **Port** – Port number. (Range: 25-28/49-52)
- ◆ **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 43: 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	3.45	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 thresholds 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 detects per pair cable faults (accuracy $\leq \pm 5$ meters).
 - Interpair shorts
 - Intrapair shorts

- Open pairs
- ◆ Detects the location of cable faults.
- ◆ Detects good cable (accuracy $\leq \pm 10$ meters).
 - Detects the length of good cables on a per pair basis
- ◆ 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 1 second. 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.
 - Pair busy, linked partner in 100BAS-TX forced mode
 - Interpair short
 - Intrapair short
 - Pair open, fault detected
 - Invalid, cable diagnostic routine did not complete successfully
- ◆ Ports must have auto-negotiation enabled
- ◆ Ports are linked down while running cable diagnostics.

Parameters

These parameters are displayed:

- ◆ **Port** – Switch port identifier.
- ◆ **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.

For link-down ports, the reported distance to a fault is accurate to within +/- 5 meters. For link-up ports, the accuracy is +/- 10 meters.

To ensure more accurate measurement of the length to a fault, first disable power-saving mode on the link partner before running cable diagnostics.
- ◆ **Accuracy** – Accuracy of distance to reported fault.
- ◆ **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 44: Performing Cable Tests

Interface > Port > Cable Test

Cable Test Port List Total: 28

Port	Test Result (Cable/Fault Distance in Meters)				Accuracy(Meters)	Last Updated	Action
	Pair A	Pair B	Pair C	Pair D			
1	OK (2)	OK (0)	Short (1)	Short (1)	10	2015-12-31 07:59:57	Test
2	Not Tested Yet	Not Tested Yet	Not Tested Yet	Not Tested Yet	0		Test
3	Not Tested Yet	Not Tested Yet	Not Tested Yet	Not Tested Yet	0		Test
4	Not Tested Yet	Not Tested Yet	Not Tested Yet	Not Tested Yet	0		Test
5	Not Tested Yet	Not Tested Yet	Not Tested Yet	Not Tested Yet	0		Test
6	Not Tested Yet	Not Tested Yet	Not Tested Yet	Not Tested Yet	0		Test
7	Not Tested Yet	Not Tested Yet	Not Tested Yet	Not Tested Yet	0		Test
8	Not Tested Yet	Not Tested Yet	Not Tested Yet	Not Tested Yet	0		Test
9	Not Tested Yet	Not Tested Yet	Not Tested Yet	Not Tested Yet	0		Test
10	Not Tested Yet	Not Tested Yet	Not Tested Yet	Not Tested Yet	0		Test

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 26 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 26 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.
- ◆ When using an LAG composed of all 10G ports, different transceiver types may be used as long as the speed of each member port is the same.
- ◆ Any of the ports on the front panel can be trunked together, including ports of different media types as long as the speed of each member port is the same.
- ◆ All the ports in a trunk have to be treated as a whole when moved from/to, added or deleted from a VLAN.
- ◆ STP, VLAN, and IGMP settings can only be made for the entire trunk.

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

Figure 45: Configuring Static Trunks



Command Usage

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

Parameters

These parameters are displayed:

- ◆ **Trunk ID** – Trunk identifier. (Range: 1-26)
- ◆ **Member** – The initial trunk member. Use the Add Member page to configure additional members.
 - **Unit** – Unit identifier. (Range: 1)
 - **Port** – Port identifier. (Range: 1- 28/52)

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 46: Creating Static Trunks

The screenshot shows a web interface for configuring static trunks. At the top, the breadcrumb navigation reads "Interface > Trunk > Static". Below this, there are two dropdown menus: "Step:" with "1. Configure Trunk" selected, and "Action:" with "Add" selected. The main configuration area contains a "Trunk ID (1-26)" label next to a text input field containing the number "1". Below this, the "Member" section has "Unit" and "Port" labels, each followed by a dropdown menu both containing the number "1". At the bottom right of the form are two buttons: "Apply" and "Revert".

To add member ports to a static trunk:

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

Figure 47: Adding Static Trunks Members

Interface > Trunk > Static

Step: 1. Configure Trunk Action: Add Member

Trunk 1

Member Unit 1 Port 1

Apply Revert

To configure connection parameters for a static trunk:

1. Click Interface, Trunk, Static.
2. Select Configure General from the Step list.
3. Select Configure from the Action list.
4. Modify the required interface settings. (Refer to [“Configuring by Port List”](#) on page 106 for a description of the parameters.)
5. Click Apply.

Figure 48: Configuring Connection Parameters for a Static Trunk

Interface > Trunk > Static

Step: 2. Configure General Action: Configure

Static Trunk List Total: 1

Trunk	Type	Name	Admin	Autonegotiation	Speed Duplex	Flow Control	MTU Size (1500-9216)	Link Up Down Trap
1	1000BASE-T		☒ Enabled	☒ Enabled 10h ☒ 100h ☒ 1000f 10f ☒ 100f ☐ Sym	100full	☒ Enabled	1518	☒ 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 49: Showing Information for Static Trunks

Interface > Trunk > Static

Step: 2. Configure General


Action: Show Information

Static Trunk List Total: 1

Trunk	Type	Name	Admin	Oper Status	Shutdown Reason	Autonegotiation	Oper Speed Duplex	Oper Flow Control	MTU Size	Link Up Down Trap
1	1000BASE-T		Enabled	Up		Enabled	100full	None	1518	Enabled

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

Figure 50: Configuring Dynamic Trunks



Command Usage

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

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



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

Parameters

These parameters are displayed:

Configure Aggregator

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

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

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

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

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

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

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

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

Configure Aggregation Port - General

- ◆ **Port** – Port identifier. (Range: 1-28/52)
- ◆ **LACP Status** – Enables or disables LACP on a port.

Configure Aggregation Port - Actor/Partner

- ◆ **Port** – Port number. (Range: 1-28/52)
- ◆ **Admin Key** – The LACP administration key must be set to the same value for ports that belong to the same LAG. (Range: 0-65535; Default – Actor: 1, Partner: 0)

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



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

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

- ◆ **System Priority** – LACP system priority is used to determine link aggregation group (LAG) membership, and to identify this device to other switches during LAG negotiations. (Range: 0-65535; Default: 32768)

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



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


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

Web Interface

To configure the admin key for a dynamic trunk:

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

Figure 51: Configuring the LACP Aggregator Admin Key



Interface > Trunk > Dynamic				
Step: 1. Configure Aggregator				
Trunk List Total: 8				
Trunk	Admin Key (0-65535)	Timeout Mode	System Priority	System MAC Address
1	0	Long Timeout	32768	00-E0-0C-00-00-FD
2	0	Long Timeout	32768	00-E0-0C-00-00-FD
3	0	Long Timeout	32768	00-E0-0C-00-00-FD
4	0	Long Timeout	32768	00-E0-0C-00-00-FD
5	0	Long Timeout	32768	00-E0-0C-00-00-FD

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 52: Enabling LACP on a Port

Interface > Trunk > Dynamic

Step: 2. Configure Aggregation Port Action: Configure

General Actor Partner

Port List Total: 28

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

To configure LACP parameters for group members:

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

Figure 53: Configuring LACP Parameters on a Port

Interface > Trunk > Dynamic

Step: 2. Configure Aggregation Port Action: Configure

General Actor Partner

Port List Total: 28

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

To show the active members of a dynamic trunk:

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

Figure 54: Showing Members of a Dynamic Trunk

Interface > Trunk > Dynamic

Step: 3. Configure Trunk Action: Show Member

Trunk: 1

Member List Total: 2

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

To configure connection parameters for a dynamic trunk:

1. Click Interface, Trunk, Dynamic.
2. Select Configure Trunk from the Step list.
3. Select Configure from the Action list.
4. Modify the required interface settings. (See [“Configuring by Port List” on page 106](#) for a description of the interface settings.)
5. Click Apply.

Figure 55: Configuring Connection Settings for a Dynamic Trunk

Interface > Trunk > Dynamic

Step: 3. Configure Trunk Action: Configure

Dynamic Trunk List Total: 1


Trunk	Type	Name	Admin	Autonegotiation	Speed Duplex	Flow Control	MTU Size (1500-9216)	Link Up Down Trap
1	1000BASE-T		Enabled	☒ Enabled ☒ 10h ☒ 100h ☒ 1000f ☐ FC ☒ 10f ☒ 100f ☐ Sym	100full	Enabled	1518	Enabled

Apply Revert

To display connection parameters for a dynamic trunk:

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

Figure 56: Displaying Connection Parameters for Dynamic Trunks



The screenshot shows a web interface with a breadcrumb trail 'Interface > Trunk > Dynamic'. Below this, there are two dropdown menus: 'Step: 3. Configure Trunk' and 'Action: Show'. The main content area is titled 'Dynamic Trunk List Total: 1'. It contains a table with the following data:

Trunk	Type	Name	Admin	Oper Status	Shutdown Reason	Autonegotiation	Oper Speed Duplex	Oper Flow Control	MTU Size	Link Up Down Trap
1	1000BASE-T		Enabled	Up		Enabled	100full	None	1518	Enabled

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

Parameters

These parameters are displayed:

Table 7: LACP Port Counters

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

Web Interface

To display LACP port counters:

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

4. Click Counters.
5. Select a group member from the Port list.

Figure 57: Displaying LACP Port Counters

Interface > Trunk > Dynamic

Step: 2. Configure Aggregation Port Action: Show Information

☒ Counters ☐ Internal ☐ Neighbors

Port 1

Trunk ID 2

Port Counters Information

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

Refresh

Displaying LACP Settings and Status for the Local Side

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

Parameters

These parameters are displayed:

Table 8: LACP Internal Configuration Information

Parameter	Description
LACP System Priority	LACP system priority assigned to this port channel.
LACP Port Priority	LACP port priority assigned to this interface within the channel group.
Admin Key	Current administrative value of the key for the aggregation port.
Oper Key	Current operational value of the key for the aggregation port.
LACPDUs Interval	Number of seconds before invalidating received LACPDU information.
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.

Table 8: LACP Internal Configuration Information (Continued)

Parameter	Description
	♦ Synchronization – The System considers this link to be IN_SYNC; i.e., it has been allocated to the correct Link Aggregation Group, the group has been associated with a compatible Aggregator, and the identity of the Link Aggregation Group is consistent with the System ID and operational Key information transmitted. ♦ Aggregation – The system considers this link to be aggregatable; i.e., a potential candidate for aggregation. ♦ Long timeout – Periodic transmission of LACPDUs uses a slow transmission rate. ♦ LACP-Activity – Activity control value with regard to this link. (0: Passive; 1: Active)

Web Interface

To display LACP settings and status for the local side:

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

Figure 58: Displaying LACP Port Internal Information

Interface > Trunk > Dynamic

Step: 2. Configure Aggregation Port Action: Show Information

☐ Counters ☒ Internal ☐ Neighbors

Port 1

Trunk ID 2

Port Internal Information

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

Displaying LACP Settings and Status for the Remote Side

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

Parameters

These parameters are displayed:

Table 9: LACP Remote Device Configuration Information

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

Web Interface

To display LACP settings and status for the remote side:

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

Figure 59: Displaying LACP Port Remote Information

Interface > Trunk > Dynamic

Step: 2. Configure Aggregation Port Action: Show Information

☐ Counters
 ☐ Internal
 ☒ Neighbors

Port 3

Trunk ID 2

Port Neighbors Information

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

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

Command Usage

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

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

Parameters

These parameters are displayed for the load balance mode:

- ◆ **Destination IP Address** - Load balancing based on destination IP address.
- ◆ **Destination MAC Address** - Load balancing based on destination MAC address.
- ◆ **Source and Destination IP Address** - Load balancing based on source and destination IP address.
- ◆ **Source and Destination MAC Address** - Load balancing based on source and destination MAC address.
- ◆ **Source IP Address** - Load balancing based on source IP address.
- ◆ **Source MAC Address** - Load balancing based on source MAC address.

Web Interface

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

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

Figure 60: Configuring Load Balancing



Saving Power

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

Command Usage

- ◆ IEEE 802.3 defines the Ethernet standard and subsequent power requirements based on cable connections operating at 100 meters. Enabling power saving mode can reduce power used for cable lengths of 60 meters or less, with more significant reduction for cables of 20 meters or less, and continue to ensure signal integrity.
- ◆ The power-saving methods provided by this switch include:
 - Power saving when there is no link partner:

Under normal operation, the switch continuously auto-negotiates to find a link partner, keeping the MAC interface powered up even if no link connection exists. When using power-savings mode, the switch checks for energy on the circuit to determine if there is a link partner. If none is detected, the switch automatically turns off the transmitter, and most of the receive circuitry (entering Sleep Mode). In this mode, the low-power energy-detection circuit continuously checks for energy on the cable. If none is detected, the MAC interface is also powered down to save additional energy. If energy is detected, the switch immediately turns on both the transmitter and receiver functions, and powers up the MAC interface.
 - Power saving when there is a link partner:

Traditional Ethernet connections typically operate with enough power to support at least 100 meters of cable even though average network cable length is shorter. When cable length is shorter, power consumption can be reduced since signal attenuation is proportional to cable length. When power-savings mode is enabled, the switch analyzes cable length to determine whether or not it can reduce the signal amplitude used on a particular link.



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

Parameters

These parameters are displayed:

- ◆ **Port** – Power saving mode only applies to the Gigabit Ethernet ports using copper media. (Range: 1-24/48)

- ◆ **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: Disabled 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 61: Enabling Power Savings


The screenshot shows a web interface titled "Interface > Green Ethernet". Below the title is a "Port Green Ethernet List" with a "Total: 28" count. The list is a table with two columns: "Port" and "Power Saving Status". The "Port" column lists ports 1 through 10. The "Power Saving Status" column shows a checkbox labeled "Enabled" for each port. At the bottom of the table are "Apply" and "Revert" buttons.

Port	Power Saving Status
1	☐ Enabled
2	☐ Enabled
3	☐ Enabled
4	☐ Enabled
5	☐ Enabled
6	☐ Enabled
7	☐ Enabled
8	☐ Enabled
9	☐ Enabled
10	☐ Enabled

Configuring Local Port Mirroring

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

Figure 62: Configuring Local Port Mirroring



Command Usage

- ◆ Traffic can be mirrored from one or more source ports to a destination port on the same switch (local port mirroring as described in this section), or from one or more source ports on remote switches to a destination port on this switch

(remote port mirroring as described in [“Configuring Remote Port Mirroring” on page 144](#)).

- ◆ Monitor port speed should match or exceed source port speed, otherwise traffic may be dropped from the monitor port.
- ◆ When mirroring port traffic, the target port must be included in the same VLAN as the source port when using MSTP (see [“Spanning Tree Algorithm” on page 205](#)).
- ◆ The destination port cannot be a trunk or trunk member port.
- ◆ Note that Spanning Tree BPDU packets are not mirrored to the target port.

Parameters

These parameters are displayed:

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

Web Interface

To configure a local mirror session:

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

Figure 63: Configuring Local Port Mirroring

The screenshot shows a web interface titled "Interface > Mirror". At the top, there is an "Action:" dropdown menu with "Show" selected. Below this, there are two rows of configuration fields. The first row is for "Source Port", with "Unit" and "Port" dropdown menus, both set to "1". The second row is for "Target Port", also with "Unit" and "Port" dropdown menus, both set to "1". Below these, there is a "Type" dropdown menu set to "Both". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

To display the configured mirror sessions:

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


Figure 64: Displaying Local Port Mirror Sessions

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

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 65: 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 163) 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 port⁶. 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 port⁷, whether or not the traffic exiting this port will be tagged or untagged, and the RSPAN VLAN. Then specify each uplink port where the mirrored traffic is being received.

◆ *RSPAN Limitations*

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

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

6. Only 802.1Q trunk ports or hybrid ports can be configured as an RSPAN uplink port (see [“Adding Static Members to VLANs”](#) on page 165).

7. If RSPAN packets are configured to carry a VLAN tag, then the destination port type cannot be access mode (see [“Adding Static Members to VLANs”](#) on page 165). If the RSPAN packets are configured as untagged, then the destination port type cannot be trunk mode.

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)
- ◆ **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 163](#)).
- ◆ **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 ports⁶ configured on an intermediate or destination switch.

Only destination and uplink ports will be assigned by the switch as members of the RSPAN VLAN. *Ports cannot be manually assigned to an RSPAN VLAN through the VLAN > Static page.* Nor can GVRP dynamically add port members to an RSPAN VLAN. Also, note that the VLAN > Static (Show) page will not display any members for an RSPAN VLAN, but will only show configured RSPAN VLAN identifiers.

- ◆ **Source Port** – Specifies the ports to monitor. (Switch Role as Source)

- ◆ **Type** – Specifies the traffic type to be mirrored remotely. (Options: Rx, Tx, Both)
- ◆ **Destination Port** – Specifies the destination port to monitor the traffic mirrored from the source ports. Only one destination port can be configured on the same switch per session, but a destination port can be configured on more than one switch for the same session. Also note that a destination port can still send and receive switched traffic, and participate in any Layer 2 protocols to which it has been assigned. (Switch Role as Intermediate or Destination)
- ◆ **Tag** – Specifies whether or not the traffic exiting the destination port to the monitoring device carries the RSPAN VLAN tag.

Web Interface

To configure a remote mirror session:

1. Click Interface, RSPAN.
2. Set the Switch Role to None, Source, Intermediate, or Destination.
3. Configure the required settings for each switch participating in the RSPAN VLAN.
4. Click Apply.

Figure 66: Configuring Remote Port Mirroring (Source)

Interface > RSPAN

Session: 1

Operation Status: Up

Switch Role: Source

Remote VLAN: 2

Uplink Port: 4

Source Port Configuration List Total: 28

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

Figure 67: Configuring Remote Port Mirroring (Intermediate)

Interface > RSPAN

Session: 1

Operation Status: Up

Switch Role: Intermediate

Remote VLAN: 2

Uplink Port List Total: 28

Port	Uplink
1	☐
2	☐
3	☐
4	☒
5	☐

Figure 68: Configuring Remote Port Mirroring (Destination)

Interface > RSPAN

Session: 1

Operation Status: Up

Switch Role: Destination

Destination Port: 1

Tag: Untagged

Remote VLAN: 2

Uplink Port List Total: 28

Port	Uplink
1	☐
2	☐
3	☐
4	☒
5	☐

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**⁸ – 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: 0-10000000 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**⁸ – 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

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

used to indicate the appropriate number of zeros required to fill the undefined fields.

- ◆ **Receiver Socket Port**⁸ – The UDP port on which the sFlow Collector is listening for sFlow streams. (Range: 1-65534)
- ◆ **Maximum Datagram Size** – Maximum size of the sFlow datagram payload. (Range: 200-1500 bytes)
- ◆ **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 69: Configuring an sFlow Receiver

The screenshot shows a web interface titled "Interface > sFlow". At the top, there are two dropdown menus: "Step: 1. Configure Receiver" and "Action: Add". Below these are several input fields for configuring the receiver:

- Receiver Owner Name:** A text box containing "stat_server1".
- Receiver Timeout (30 - 10000000):** A text box containing "100" followed by a "sec" label.
- Receiver Destination:** A text box containing "192.168.220.225".
- Receiver Socket Port (1 - 65535):** A text box containing "22500".
- Maximum Datagram Size (200 - 1500):** A text box containing "512" followed by a "bytes" label.
- Datagram Version:** A dropdown menu with a checked box next to "v5".


At the bottom right of the form are two buttons: "Apply" and "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 70: Showing sFlow Receivers



The screenshot shows a web interface titled "Interface > sFlow". It has a "Step:" dropdown set to "1. Configure Receiver" and an "Action:" dropdown set to "Show". Below this is a "Receiver List" with a "Total: 1" count. The table has columns: Owner Name, Timeout, Destination, Socket Port, Maximum Datagram Size, and Datagram Version. One entry is shown with Owner Name "stat_server1", Timeout "97", Destination "192.168.220.225", Socket Port "22500", Maximum Datagram Size "512", and Datagram Version "5". There are "Delete" and "Revert" buttons at the bottom right of the table.

	Owner Name	Timeout	Destination	Socket Port	Maximum Datagram Size	Datagram Version
☐	stat_server1	97	192.168.220.225	22500	512	5

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 71: Configuring an sFlow Instance

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

- Receiver Owner Name: stat_server1
- Type: Sampling (selected), Polling
- Data Source: Unit 1, Port 3
- Instance ID (1-1): 1
- Sampling Rate (256-16777215): 256
- Maximum Header Size (64-256): 200 bytes

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

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 72: Showing sFlow Instances

The screenshot shows the 'Interface > sFlow' configuration page with the 'Action' set to 'Show'. The configuration fields are the same as in Figure 71. Below the configuration fields, there is a 'Sampling List' with a total of 1 instance.

	Data Source (Unit/Port)	Instance ID	Rate	Maximum Header Size (bytes)
☐	1/3	1	256	200

Buttons for 'Delete' and 'Revert' are at the bottom right.

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 73: Enabling Traffic Segmentation

Interface > Traffic Segmentation

Step: 1. Configure Global

Status ☐ Enabled

Uplink-to-Uplink Mode: Blocking

Apply Revert

Configuring Uplink and Downlink Ports

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

Command Usage

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

Table 10: Traffic Segmentation Forwarding

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

* The forwarding state for uplink-to-uplink ports is configured on the Configure Global page (see [page 153](#)).

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

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

Parameters

These parameters are displayed:

- ◆ **Session ID** – Traffic segmentation session. (Range: 1-4)
- ◆ **Direction** – Adds an interface to the segmented group by setting the direction to uplink or downlink. (Default: Uplink)
- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Port** – Port Identifier. (Range: 1-28/52)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-26)

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 74: Configuring Members for Traffic Segmentation

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

To show the members of the traffic segmentation group:

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

3. Select Show from the Action list.

Figure 75: Showing Traffic Segmentation Members

Interface > Traffic Segmentation

Step: 2. Configure Session Action: Show

Session List Total: 2

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

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 165](#)). If VLAN trunking is enabled on an interface, then that interface cannot be set to access mode, and vice versa.

- ◆ To prevent loops from forming in the spanning tree, all unknown VLANs will be bound to a single instance (either STP/RSTP or an MSTP instance, depending on the selected STA mode).
- ◆ If both VLAN trunking and ingress filtering are disabled on an interface, packets with unknown VLAN tags will still be allowed to enter this interface and will be flooded to all other ports where VLAN trunking is enabled. (In other words, VLAN trunking will still be effectively enabled for the unknown VLAN).

Parameters

These parameters are displayed:

- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Port** – Port Identifier. (Range: 1-28/52)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-26)
- ◆ **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: 28

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

Apply Revert

VLAN Configuration

This chapter includes the following topics:

- ◆ [IEEE 802.1Q VLANs](#) – Configures static and dynamic VLANs.
- ◆ [IEEE 802.1Q Tunneling](#) – Configures QinQ tunneling to maintain customer-specific VLAN and Layer 2 protocol configurations across a service provider network, even when different customers use the same internal VLAN IDs.
- ◆ [L2PT Tunneling](#) – Configures Layer 2 Protocol Tunneling for the specified protocol.
- ◆ [Protocol VLANs](#)⁹ – Configures VLAN groups based on specified protocols.
- ◆ [IP Subnet VLANs](#)⁹ – 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](#)⁹ – 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 Translation](#) – Maps VLAN IDs between the customer and the service provider.

IEEE 802.1Q VLANs

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

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

VLANs help to simplify network management by allowing you to move devices to a new VLAN without having to change any physical connections. VLANs can be easily organized to reflect departmental groups (such as Marketing or R&D), usage

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

groups (such as e-mail), or multicast groups (used for multimedia applications such as video conferencing).

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

This switch supports the following VLAN features:

- ◆ Up to 4094 VLANs based on the IEEE 802.1Q standard
- ◆ Distributed VLAN learning across multiple switches using explicit or implicit tagging and GVRP protocol
- ◆ Port overlapping, allowing a port to participate in multiple VLANs
- ◆ End stations can belong to multiple VLANs
- ◆ Passing traffic between VLAN-aware and VLAN-unaware devices
- ◆ Priority tagging


Assigning Ports to VLANs

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



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

Figure 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 165](#)). 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.
VLAN 4093 is dedicated for [Switch Clustering](#). Configuring this VLAN for other purposes may cause problems in the Clustering operation.
- ◆ **Status** – Enables or disables the specified VLAN.
- ◆ **Remote VLAN** – Reserves this VLAN for RSPAN (see [“Configuring Remote Port Mirroring” on page 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 639](#)).

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. Enable the Status field to set 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 a web interface for configuring static VLANs. The title bar reads "VLAN > Static". Below the title bar, there is a section labeled "Action:" with a dropdown menu set to "Add". Underneath, there are three rows of configuration fields: "VLAN ID (1-4094)" with a text input containing "2" and a range selector "-"; "Status" with a checked checkbox and the label "Enabled"; and "Remote VLAN" with an unchecked checkbox and the label "Enabled". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

To modify the configuration settings for VLAN groups:

1. Click VLAN, Static.
2. Select Modify from the Action list.
3. Select the identifier of a configured VLAN.
4. Modify the VLAN name 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

VLAN > Static

Action:

Modify

VLAN ID (1-4093)

2

VLAN Name

R&D

Status

☒ Enabled

L3 Interface

☒ Enabled

Apply

Revert

To show the configuration settings for VLAN groups:

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

Figure 82: Showing Static VLANs

VLAN > Static

Action:

Show

Static VLAN List

Total: 2

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

Delete

Revert

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: 1-28/52)
 - **Trunk** – Trunk Identifier. (Range: 1-26)
- ◆ **Mode** – Indicates VLAN membership mode for an interface. (Default: Hybrid)
 - **Access** – Sets the port to operate as an untagged interface. The port transmits and receives untagged frames on a single VLAN only.

Access mode is mutually exclusive with VLAN trunking (see [“VLAN Trunking” on page 156](#)). If VLAN trunking is enabled on an interface, then that interface cannot be set to access mode, and vice versa.
 - **Hybrid** – Specifies a hybrid VLAN interface. The port may transmit tagged or untagged frames.
 - **1Q Trunk** – Specifies a port as an end-point for a VLAN trunk. A trunk is a direct link between two switches, so the port transmits tagged frames that identify the source VLAN. Note that frames belonging to the port’s default VLAN (i.e., associated with the PVID) are also transmitted as tagged frames.
- ◆ **PVID** – VLAN ID assigned to untagged frames received on the interface. (Default: 1)

When using Access mode, and an interface is assigned to a new VLAN, its PVID is automatically set to the identifier for that VLAN. When using Hybrid mode, the PVID for an interface can be set to any VLAN for which it is an untagged member.
- ◆ **Acceptable Frame Type** – Sets the interface to accept all frame types, including tagged or untagged frames, or only tagged frames. When set to receive all frame types, any received frames that are untagged are assigned to the default VLAN. (Options: All, Tagged; Default: All)
- ◆ **Ingress Filtering** – Determines how to process frames tagged for VLANs for which the ingress port is not a member. (Default: Disabled)
 - Ingress filtering only affects tagged frames.
 - If ingress filtering is disabled and a port receives frames tagged for VLANs for which it is not a member, these frames will be flooded to all other ports (except for those VLANs explicitly forbidden on this port).
 - If ingress filtering is enabled and a port receives frames tagged for VLANs for which it is not a member, these frames will be discarded.
 - Ingress filtering does not affect VLAN independent BPDU frames, such as GVRP or STP. However, they do affect VLAN dependent BPDU frames, such as GMRP.

- ◆ **Membership Type** – Select VLAN membership for each interface by marking the appropriate radio button for a port or trunk:
 - **Tagged:** Interface is a member of the VLAN. All packets transmitted by the port will be tagged, that is, carry a tag and therefore carry VLAN or CoS information.
 - **Untagged:** Interface is a member of the VLAN. All packets transmitted by the port will be untagged, that is, not carry a tag and therefore not carry VLAN or CoS information. Note that an interface must be assigned to at least one group as an untagged port.
 - **Forbidden:** Interface is forbidden from automatically joining the VLAN via GVRP. For more information, see “Automatic VLAN Registration” on page 161. This attribute cannot be configured if the specified VLAN does not exist on the switch.
 - **None:** Interface is not a member of the VLAN. Packets associated with this VLAN will not be transmitted by the interface.



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

Edit Member by Interface

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

Edit Member by Interface Range

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

- ◆ **Port Range** – Displays a list of ports. (Range: 1-28/52)
- ◆ **Trunk Range** – Displays a list of ports. (Range: 1-26)



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: 28 1 2 3

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	☐	☒	☐	☐
4	Hybrid	1	All	☒ Enabled	☐	☒	☐	☐
5	Hybrid	1	All	☒ Enabled	☐	☒	☐	☐

To configure static members by interface:

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

Figure 84: Configuring Static VLAN Members by Interface

VLAN > Static

Action: Edit Member by Interface

Interface: ☒ Port 1 ☐ Trunk

Mode: Hybrid

PVID: 1

Acceptable Frame Type: All

Ingress Filtering: ☐ Enabled

Static VLAN Membership List Total: 4

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

Apply Revert

To configure static members by interface range:

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

Figure 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 (1-4093): -

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)

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

Configure Interface

- ◆ **Interface** – Displays a list of ports or trunks.
 - **Port** – Port Identifier. (Range: 1-28/52)
 - **Trunk** – Trunk Identifier. (Range: 1-26)
- ◆ **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)
- ◆ **GARP 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

The screenshot shows a web interface for configuring GVRP. At the top, it says "VLAN > Dynamic". Below that, there is a "Step:" dropdown menu with "1. Configure General" selected. Underneath, there is a "GVRP Status" section with a checked checkbox and the word "Enabled". At the bottom right, there are two buttons: "Apply" and "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	GARP Timer (centiseconds)		
		Join (20-1000)	Leave (60-3000)	LeaveAll (500-18000)
1	☐ Enabled	20	60	1000
2	☐ Enabled	20	60	1000
3	☒ Enabled	20	60	1000
4	☐ Enabled	20	60	1000
5	☐ Enabled	20	60	1000

To show the dynamic VLAN joined by this switch:

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

Figure 88: Showing Dynamic VLANs Registered on the Switch

VLAN > Dynamic

Step: 3. Show Dynamic VLAN Action: Show VLAN

Dynamic VLAN List Total: 2

VLAN ID	VLAN Name	Status
120		Enabled
150		Enabled

To show the members of a dynamic VLAN:

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

Figure 89: Showing the Members of a Dynamic VLAN

VLAN > Dynamic

Step: 3. Show Dynamic VLAN Action: Show VLAN Member

VLAN 120

Dynamic VLAN Member List Total: 10

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

IEEE 802.1Q Tunneling

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

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


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

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

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

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

Figure 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 178](#). The priority of the inner tag is copied to the outer tag if it is a tagged or priority tagged packet.
2. After successful source and destination lookup, the ingress process sends the packet to the switching process with two tags. If the incoming packet is untagged, the outer tag is an SPVLAN tag, and the inner tag is a dummy tag (8100 0000). If the incoming packet is tagged, the outer tag is an SPVLAN tag, and the inner tag is a CVLAN tag.

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

Layer 2 Flow for Packets Coming into a Tunnel Uplink Port

An uplink port receives one of the following packets:

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

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

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

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

Configuration Limitations for QinQ

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

General Configuration Guidelines for QinQ

1. Enable Tunnel Status, and set the Tag Protocol Identifier (TPID) value of the tunnel access port (in the Ethernet Type field). This step is required if the attached client is using a nonstandard 2-byte ethertype to identify 802.1Q tagged frames. The default ethertype value is 0x8100. (See [“Enabling QinQ Tunneling on the Switch” on page 177](#).)
2. Create a Service Provider VLAN, also referred to as an SPVLAN (see [“Configuring VLAN Groups” on page 163](#)).
3. Configure the QinQ tunnel access port to Access mode (see [“Adding an Interface to QinQ Tunnel” on page 180](#)).
4. Configure the QinQ tunnel access port to join the SPVLAN as an untagged member (see [“Adding Static Members to VLANs” on page 165](#)).

5. Configure the SPVLAN ID as the native VID on the QinQ tunnel access port (see [“Adding Static Members to VLANs” on page 165](#)).
6. Configure the QinQ tunnel uplink port to Uplink mode (see [“Adding an Interface to QinQ Tunnel” on page 180](#)).
7. Configure the QinQ tunnel uplink port to join the SPVLAN as a tagged member (see [“Adding Static Members to VLANs” on page 165](#)).

Enabling QinQ Tunneling on the Switch

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

Parameters

These parameters are displayed:

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

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

The specified ethertype only applies to ports configured in Uplink mode (see [“Adding an Interface to QinQ Tunnel” on page 180](#)). 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

VLAN > Tunnel

Step: 1. Configure Global

Tunnel Status ☒ Enabled

Ethernet Type
(8000-FFFF, hexadecimal value)

Apply Revert

Creating CVLAN to SPVLAN Mapping Entries

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

Command Usage

- ◆ The inner VLAN tag of a customer packet entering the edge router of a service provider's network is mapped to an outer tag indicating the service provider VLAN that will carry this traffic across the 802.1Q tunnel. By default, the outer tag is based on the default VID of the edge router's ingress port. This process is performed in a transparent manner as described under ["IEEE 802.1Q Tunneling" on page 173](#).
- ◆ When priority bits are found in the inner tag, these are also copied to the outer tag. This allows the service provider to differentiate service based on the indicated priority and appropriate methods of queue management at intermediate nodes across the tunnel.
- ◆ Rather than relying on standard service paths and priority queuing, QinQ VLAN mapping can be used to further enhance service by defining a set of differentiated service pathways to follow across the service provider's network for traffic arriving from specified inbound customer VLANs.
- ◆ Note that all customer interfaces should be configured as access interfaces (that is, a user-to-network interface) and service provider interfaces as uplink interfaces (that is, a network-to-network interface). Use the Configure Interface page described in the next section to set an interface to access or uplink mode.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28/52)
- ◆ **Customer VLAN ID** – VLAN ID for the inner VLAN tag. (Range: 1-4094)

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

Web Interface

To configure a mapping entry:

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

Figure 92: Configuring CVLAN to SPVLAN Mapping Entries

VLAN > Tunnel

Step: 2. Configure Service Action: Add

Port: 1

Customer VLAN ID (1-4094):

Service VLAN ID (1-4093):

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

Port: 1

Tunnel Service Subscriptions List Total: 2

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

Delete Revert

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

Adding an Interface to QinQ Tunnel

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

Command Usage

- ◆ Use the Configure Global page to set the switch to QinQ mode before configuring a tunnel access port or tunnel uplink port (see [“Enabling QinQ Tunneling on the Switch” on page 177](#)). 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: 1-28/52)
 - **Trunk** – Trunk Identifier. (Range: 1-26)
- ◆ **Mode** – Sets the VLAN membership mode of the port.
 - **None** – The port operates in its normal VLAN mode. (This is the default.)
 - **Access** – Configures QinQ tunneling for a client access port to segregate and preserve customer VLAN IDs for traffic crossing the service provider network.
 - **Uplink** – Configures QinQ tunneling for an uplink port to another device within the service provider network.

Web Interface

To add an interface to a QinQ tunnel:

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

Figure 94: Adding an Interface to a QinQ Tunnel

VLAN > Tunnel

Step: 2. Configure Interface

Interface ☒ Port ☐ Trunk

802.1Q Tunnel Port List Total: 52

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

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 ports¹⁰ 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 access port, or (3) received packet is Cisco-compatible L2PT (i.e., as indicated by a proprietary MAC address).

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

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 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. At the top, it says "VLAN > L2PT". Below that is a "Step:" dropdown menu with "1. Configure Global" selected. The main section is titled "Tunnel MAC Address" and contains a text input field with the value "01-12-CF-00-00-02". To the right of the input field is a placeholder text "(XX-XX-XX-XX-XX-XX of 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
- ◆ **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

VLAN > L2PT					
Step: 2. Configure Interface ▼					
L2PT Interface List Total: 52					
Interface	CDP	LLDP	PVST+	Spanning Tree	VTP
Eth1/1	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled
Eth1/2	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled
Eth1/3	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled
Eth1/4	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled	☐ Enabled
Eth1/5	☐ 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 ([“Configuring VLAN Groups” on page 163](#)). Although not mandatory, we suggest configuring a separate VLAN for each major protocol running on your network. Do not add port members at this time.
 2. Create a protocol group for each of the protocols you want to assign to a VLAN using the Configure Protocol (Add) page.
 3. Then map the protocol for each interface to the appropriate VLAN using the Configure Interface (Add) page.
- ◆ When MAC-based, IP subnet-based, or protocol-based VLANs are supported concurrently, priority is applied in this sequence, and then port-based VLANs last.

Configuring Protocol VLAN Groups

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

Parameters

These parameters are displayed:

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



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

Web Interface

To configure a protocol group:

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

Figure 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	80 35	2
☐	RFC 1042	08 00	1
☐	RFC 1042	80 35	3
☐	LLC Other	FF FF	5

Delete Revert

Mapping Protocol Groups to Interfaces

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

Command Usage

- ◆ When creating a protocol-based VLAN, only assign interfaces using this configuration screen. If you assign interfaces using any of the other VLAN menus such as the VLAN Static table (page 165), these interfaces will admit traffic of any protocol type into the associated VLAN.
- ◆ When a frame enters a port that has been assigned to a protocol VLAN, it is processed in the following manner:
 - If the frame is tagged, it will be processed according to the standard rules applied to tagged frames.

- If the frame is untagged and the protocol type matches, the frame is forwarded to the appropriate VLAN.
- If the frame is untagged but the protocol type does not match, the frame is forwarded to the default VLAN for this interface.

Parameters

These parameters are displayed:

- ◆ **Interface** – Displays a list of ports or trunks.
 - **Port** – Port Identifier. (Range: 1-28/52)
 - **Trunk** – Trunk Identifier. (Range: 1-26)
- ◆ **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)

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. Click Apply.

Figure 99: Assigning Interfaces to Protocol VLANs

The screenshot shows a web interface for configuring Protocol VLANs. The breadcrumb is 'VLAN > Protocol'. The 'Step' dropdown is set to '2. Configure Interface' and the 'Action' dropdown is set to 'Add'. The configuration area has three rows: 'Interface' with radio buttons for 'Port' (selected) and 'Trunk', and a dropdown for '1'; 'Protocol Group ID' with a dropdown for '1'; and 'VLAN ID (1-4094)' with an empty text input field. 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. At the top, there's a breadcrumb 'VLAN > Protocol'. Below it, a 'Step:' dropdown is set to '2. Configure Interface' and an 'Action:' dropdown is set to 'Show'. Under the 'Interface' section, 'Port' is selected with a dropdown showing '2', and 'Trunk' is unselected. Below this is the 'Port To Protocol Group Mapping Table' with a 'Total: 2' indicator. The table has two columns: 'Protocol Group ID' and 'VLAN ID'. It contains two rows: one with Protocol Group ID '1' and VLAN ID '10', and another with Protocol Group ID '3' and VLAN ID '20'. Each row has a checkbox on the left. At the bottom of the table are 'Delete' and 'Revert' buttons.

	Protocol Group ID	VLAN ID
☐	1	10
☐	3	20

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. For example, ping requests must pass through and 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:

Subnet Mask:

VLAN (1-4093):

Priority (0-7):

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 or xxxxxxxxxxxx.
- ◆ **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

The screenshot shows a web interface for configuring MAC-based VLANs. The title bar reads "VLAN > MAC-Based". Below the title bar, there is an "Action:" label followed by a dropdown menu set to "Add". The main configuration area contains four rows of fields: "MAC Address" with the value "00-ab-cd-11-22-33", "Mask" with an empty field, "VLAN (1-4093)" with the value "10", and "Priority (0-7)" with an empty field. At the bottom right of the form are two buttons: "Apply" and "Revert".

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



The screenshot shows the 'VLAN > MAC-Based' configuration page. At the top, there is a breadcrumb 'VLAN > MAC-Based' and a green help icon. Below this is an 'Action:' dropdown menu set to 'Show'. The main section is titled 'MAC-Based VLAN List' with a 'Total: 1' indicator. It contains a table with the following data:

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

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

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 105: 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. (Range: 1-28/52)
- ◆ **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 106: Configuring VLAN Translation



The screenshot shows the 'VLAN > Translation' web interface. At the top, there is a header 'VLAN > Translation'. Below it, there is a form with the following fields:

- Action:** A dropdown menu with 'Add' selected.
- Port:** A dropdown menu with '1' selected.
- Old VLAN (1-4093):** A text input field containing '10'.
- New VLAN (1-4093):** A text input field containing '100'.

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

To show the mapping entries for VLANs translation:

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

Figure 107: Showing the Entries for VLAN Translation

VLAN > Translation

Action: Show

Port 1

VLAN Translation Table Total: 1

	Old VLAN	New VLAN
☐	10	100

Delete

Revert

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 Notification Traps](#) – Issue a trap when a dynamic MAC address is added or removed.

Dynamic Address Cache

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 108: Displaying the Dynamic MAC Address Table

The screenshot shows the 'MAC Address > Dynamic' web interface. At the top, the 'Action' dropdown is set to 'Show Dynamic MAC'. Below this, the 'Query by:' section has a 'Sort Key' dropdown set to 'MAC Address'. There are three checkboxes: 'MAC Address' (unchecked), 'VLAN' (checked), and 'Interface' (unchecked). The 'VLAN' checkbox is selected, and the 'VLAN' dropdown is set to '1'. The 'Interface' checkbox is unchecked, and the 'Interface' dropdown is set to 'Port 1'. A 'Query' button is located to the right of the search parameters. Below the search parameters, the 'Dynamic MAC Address List' is displayed with a 'Total: 2' count. The table has five columns: 'MAC Address', 'VLAN', 'Interface', 'Type', and 'Life Time'. The table contains two entries: one with MAC Address '00-E0-29-94-34-64' and another with '70-72-CF-32-DD-FF', both in VLAN 1, Interface 'Unit 1 / Port 1', Type 'Learn', and Life Time 'Delete on Timeout'.

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

Clearing the Dynamic Address Table

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

Parameters

These parameters are displayed:

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

Web Interface

To clear the entries in the dynamic address table:

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

Figure 109: Clearing Entries in the Dynamic MAC Address Table

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

Changing the Aging Time

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

Parameters

These parameters are displayed:

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

Web Interface

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

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

5. Click Apply.

Figure 110: 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. Underneath, the "Aging Status" is set to "Enabled" with a checked checkbox. The "Aging Time (10-1000000)" is set to "300" seconds. At the bottom right, there 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 201](#)) 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 348](#)).
 - Security Status (see ["Configuring Port Security" on page 344](#)) is enabled on the same interface.

Parameters

These parameters are displayed:

- ◆ **Interface** – Displays a list of ports or trunks.
 - **Port** – Port Identifier. (Range: 1-28/52)
 - **Trunk** – Trunk Identifier. (Range: 1-26)


- ◆ **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 111: Configuring MAC Address Learning



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

Setting Static Addresses

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

Command Usage

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

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

Parameters

These parameters are displayed:

Add Static Address

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

Show Static Address

The following additional fields are also 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 112: 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 several fields: "VLAN" with a dropdown menu set to "1", "Interface" with radio buttons for "Port" (selected) and "Trunk", and a dropdown menu set to "1", "MAC Address" with a text input field containing "00-12-cf-94-34-da", and "Static Status" with a dropdown menu set to "Permanent". At the bottom right of the form, 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 113: 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 a green help icon. Below this is an 'Action:' dropdown menu set to 'Show'. The main content area is titled 'Static MAC Address to Interface Mapping Table' with a 'Total: 2' indicator. It contains a table with the following data:

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

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

Issuing MAC Address Traps

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

Parameters

These parameters are displayed:

Configure Global

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

Configure Interface

- ◆ **Port** – Port Identifier. (Range: 1-28/52)
- ◆ **MAC Notification Trap** – Enables MAC authentication traps on the current interface. (Default: Disabled)

MAC authentication traps must be enabled at the global level for this attribute to take effect.

Web Interface

To enable MAC address traps at the global level:

1. Click MAC Address, MAC Notification.
2. Select Configure Global from the Step list.

3. Configure MAC notification traps and the transmission interval.
4. Click Apply.

Figure 114: Issuing MAC Address Traps (Global Configuration)

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

To enable MAC address traps at the interface level:

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

Figure 115: Issuing MAC Address Traps (Interface Configuration)

The screenshot shows the 'MAC Address > MAC Notification' configuration page, Step 2: Configure Interface. The 'Step' dropdown is set to '2. Configure Interface'. There are radio buttons for 'Interface', 'Port' (selected), and 'Trunk'. Below is a 'Port List' table with 28 total ports. The table has two columns: 'Port' and 'MAC Notification Trap'. The first five rows are visible, each with a port number and an 'Enabled' checkbox.

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

Spanning Tree Algorithm

This chapter describes the following basic topics:

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

Overview

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


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

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

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

ports, and disables all other ports. Network packets are therefore only forwarded between root ports and designated ports, eliminating any possible network loops.

Figure 116: STP Root Ports and Designated Ports




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

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


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

Figure 117: 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 223](#)). 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 118: Spanning Tree – Common Internal, Common, Internal



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

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

Configuring Loopback Detection

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

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



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

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

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

Parameters

These parameters are displayed:

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

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 119: Configuring Port Loopback Detection



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

Configuring Global Settings for STA

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

Command Usage

◆ Spanning Tree Protocol¹¹

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.

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

◆ **Rapid Spanning Tree Protocol¹¹**

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

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

◆ **Multiple Spanning Tree Protocol**

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

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

Parameters

These parameters are displayed:

Basic Configuration of Global Settings

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

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

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

◆ **BPDU Flooding** – Configures the system to flood BPDUs to all other ports on the switch or just to all other ports in the same VLAN when spanning tree is disabled globally on the switch or disabled on a specific port.

- To VLAN: Floods BPDUs to all other ports within the receiving port's native VLAN (i.e., as determined by port's PVID). This is the default.
- To All: Floods BPDUs to all other ports on the switch.

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

Advanced Configuration Settings

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

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

When the Switch Becomes Root

- ◆ **Hello Time** – Interval (in seconds) at which the root device transmits a configuration message.
 - Default: 2
 - Minimum: 1
 - Maximum: The lower of 10 or $[(\text{Max. Message Age} / 2) - 1]$
- ◆ **Maximum Age** – The maximum time (in seconds) a device can wait without receiving a configuration message before attempting to reconverge. All device ports (except for designated ports) should receive configuration messages at regular intervals. Any port that ages out STA information (provided in the last configuration message) becomes the designated port for the attached LAN. If it

is a root port, a new root port is selected from among the device ports attached to the network. (References to “ports” in this section mean “interfaces,” which includes both ports and trunks.)

- Default: 20
- Minimum: The higher of 6 or $[2 \times (\text{Hello Time} + 1)]$
- Maximum: The lower of 40 or $[2 \times (\text{Forward Delay} - 1)]$

- ◆ **Forward Delay** – The maximum time (in seconds) this device will wait before changing states (i.e., discarding to learning to forwarding). This delay is required because every device must receive information about topology changes before it starts to forward frames. In addition, each port needs time to listen for conflicting information that would make it return to a discarding state; otherwise, temporary data loops might result.

- Default: 15
- Minimum: The higher of 4 or $[(\text{Max. Message Age} / 2) + 1]$
- Maximum: 30

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

Configuration Settings for MSTP

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



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

Web Interface

To configure global STA settings:

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

Figure 120: Configuring Global Settings for STA (STP)

The screenshot shows a web interface for configuring Spanning Tree settings for a Specific Tree (STA). The title bar reads "Spanning Tree > STA". Below the title bar, there are two dropdown menus: "Step:" with "1. Configure Global" selected, and "Action:" with "Configure" selected. The main configuration area is divided into several sections:

- Spanning Tree Status:** A checkbox labeled "Enabled" is checked.
- Spanning Tree Type:** A dropdown menu is set to "STP".
- Priority (0-61440, in steps of 4096):** A text input field contains the value "32768".
- BPDU Flooding:** A dropdown menu is set to "To VLAN".
- Advanced:** This section is expanded, showing:
 - Path Cost Method:** A dropdown menu is set to "Long".
 - Transmission Limit (1-10):** A text input field contains the value "3".
- When the Switch Becomes Root:** This section contains three text input fields with units in seconds:
 - Hello Time (1-10):** "2" sec
 - Maximum Age (6-40):** "20" sec
 - Forward Delay (4-30):** "15" sec

Below these fields, a note states: "Note: 2 * (Hello Time + 1) <= Max Age <= 2 * (Forward Delay - 1)". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

Figure 121: 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

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 122: 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

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 33

Configuration Digest 0xAC36177F50283CD4B83821D8AB26DE62

Region Revision (0-65535) 0

Region Name 00 00 e8 93 82 a0

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 123: Displaying Global Settings for STA

Spanning Tree > STA

Step: 1. Configure Global Action: Show Information

Spanning Tree Information

Spanning Tree Status	Enabled	Spanning Tree Type	RSTP
Designated Root	32768.0000E89382A0	Bridge ID	32768.0000E89382A0
Root Port	0	Max Age	20 sec
Root Path Cost	0	Hello Time	2 sec
Configuration Changes	2	Forward Delay	15 sec
Last Topology Change	0 days, 3 hours, 51 minutes, 11 seconds		

Configuring Interface Settings for STA

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

Parameters

These parameters are displayed:

- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Spanning Tree** – Enables/disables STA on this interface. (Default: Enabled)
- ◆ **BPDU Flooding** – Enables/disables the flooding of BPDUs to other ports when global spanning tree is disabled ([page 209](#)) or when spanning tree is disabled on specific port. When flooding is enabled, BPDUs are flooded to all other ports on the switch or to all other ports within the receiving port’s native VLAN as specified by the Spanning Tree BPDU Flooding attribute ([page 209](#)). (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 method¹², 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 124: 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.

12. Refer to “Configuring Global Settings for STA” on page 209 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 209](#)).

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

- If spanning tree mode is set to STP ([page 209](#)), edge-port mode cannot automatically transition to operational edge-port state using the automatic setting.
- If loopback detection is enabled ([page 208](#)) 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 220](#)).

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

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

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

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

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

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

format (RSTP or STP-compatible) to send on the selected interfaces.
(Default: Disabled)


- ◆ **TC Propagate Stop** – Stops the propagation of topology change notifications (TCN). (Default: Disabled)

Web Interface

To configure interface settings for STA:

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

Figure 125: Configuring Interface Settings for STA



The screenshot shows the 'Spanning Tree > STA' configuration page. At the top, there are tabs for 'Steps' (1. Spanning Tree, 2. Configure Interface) and 'Action' (Configure). Below this, there are radio buttons for 'Interface' (selected) and 'Trunk'. A 'Port List' section shows 'Total: 28' ports. The main table displays settings for 5 ports. Each row has columns for Port, Spanning Tree, BPDU Flooding, Priority, Admin Path Cost, Admin Link Type, Root Guard, Admin Edge Port, BPDU Guard, BPDU Guard Auto Recovery, BPDU Guard Auto Recovery Interval, BPDU Filter, Migration, and TC Propagate Stop. All settings are currently set to 'Enabled' or 'Auto'.

Port	Spanning Tree	BPDU Flooding	Priority (0-240, in steps of 16)	Admin Path Cost (0-200000000, 0: Auto)	Admin Link Type	Root Guard	Admin Edge Port	BPDU Guard	BPDU Guard Auto Recovery	BPDU Guard Auto Recovery Interval (30-36400)	BPDU Filter	Migration	TC Propagate Stop
1	Enabled	Enabled	128	0	Auto	Enabled	Auto	Enabled	Enabled	300	Enabled	Enabled	Enabled
2	Enabled	Enabled	128	0	Auto	Enabled	Auto	Enabled	Enabled	300	Enabled	Enabled	Enabled
3	Enabled	Enabled	128	0	Auto	Enabled	Auto	Enabled	Enabled	300	Enabled	Enabled	Enabled
4	Enabled	Enabled	128	0	Auto	Enabled	Auto	Enabled	Enabled	300	Enabled	Enabled	Enabled
5	Enabled	Enabled	128	0	Auto	Enabled	Auto	Enabled	Enabled	300	Enabled	Enabled	Enabled

Displaying Interface Settings for STA

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

Parameters

These parameters are displayed:

- ◆ **Spanning Tree** – Shows if STA has been enabled on this interface.
- ◆ **BPDU Flooding** – Shows if BPDUs will be flooded to other ports when spanning tree is disabled globally on the switch or disabled on a specific port.

◆ **STA Status** – Displays current state of this port within the Spanning Tree:

- **Discarding** - Port receives STA configuration messages, but does not forward packets.
- **Learning** - Port has transmitted configuration messages for an interval set by the Forward Delay parameter without receiving contradictory information. Port address table is cleared, and the port begins learning addresses.
- **Forwarding** - Port forwards packets, and continues learning addresses.


The rules defining port status are:

- A port on a network segment with no other STA compliant bridging device is always forwarding.
 - If two ports of a switch are connected to the same segment and there is no other STA device attached to this segment, the port with the smaller ID forwards packets and the other is discarding.
 - All ports are discarding when the switch is booted, then some of them change state to learning, and then to forwarding.
- ◆ **Forward Transitions** – The number of times this port has transitioned from the Learning state to the Forwarding state.
- ◆ **Designated Cost** – The cost for a packet to travel from this port to the root in the current Spanning Tree configuration. The slower the media, the higher the cost.
- ◆ **Designated Bridge** – The bridge priority and MAC address of the device through which this port must communicate to reach the root of the Spanning Tree.
- ◆ **Designated Port** – The port priority and number of the port on the designated bridging device through which this switch must communicate with the root of the Spanning Tree.
- ◆ **Oper Path Cost** – The contribution of this port to the path cost of paths towards the spanning tree root which include this port.
- ◆ **Oper Link Type** – The operational point-to-point status of the LAN segment attached to this interface. This parameter is determined by manual configuration or by auto-detection, as described for Admin Link Type in STA Port Configuration on [page 216](#).
- ◆ **Oper Edge Port** – This parameter is initialized to the setting for Admin Edge Port in STA Port Configuration on [page 216](#) (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 126: 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 127: Displaying Interface Settings for STA

Spanning Tree > STA

Step: 2. Configure Interface Action: Show Information

Interface ☒ Port ☐ Trunk

Spanning Tree Port List Total: 28

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

Configuring Multiple Spanning Trees

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

Command Usage

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

By default all VLANs are assigned to the Internal Spanning Tree (MST Instance 0) that connects all bridges and LANs within the MST region. This switch supports up to 33 instances. You should try to group VLANs which cover the same general area of your network. However, remember that you must configure all bridges within the same MSTI Region (page 209) 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 209).
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 128: Creating an MST Instance

Spanning Tree > MSTP

Step: 1. Configure Global Action: Add

MST ID (0-4094) 1

VLAN ID (1-4094) 1

Priority (0-61440, in steps of 4096)

Apply 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 129: Displaying MST Instances

The screenshot shows the 'Spanning Tree > MSTP' configuration page. At the top, there is a 'Step' dropdown set to '1. Configure Global' and an 'Action' dropdown set to 'Show'. Below this is a table titled 'MST List' with a 'Total: 2' count. The table has two columns: a checkbox column and an 'MST ID' column. There are two rows, one for MST ID 0 and one for MST ID 1. At the bottom right of the table are 'Delete' and 'Revert' buttons.

	MST ID
☐	0
☐	1

To modify the priority for an MST instance:

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

Figure 130: Modifying the Priority for an MST Instance

The screenshot shows the 'Spanning Tree > MSTP' configuration page with the 'Action' dropdown set to 'Modify'. Below the 'MST List' is a table titled 'MST Details List' with a 'Total: 2' count. The table has two columns: 'MST ID' and 'Priority (0-61440, in steps of 4096)'. There are two rows, one for MST ID 0 and one for MST ID 1. The priority for MST ID 0 is 0, and for MST ID 1 it is 32768. At the bottom right of the table are 'Apply' and 'Revert' buttons.

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

To display global settings for MSTP:

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

4. Select an MST ID. The attributes displayed on this page are described under “Displaying Global Settings for STA” on page 215.

Figure 131: Displaying Global Settings for an MST Instance

The screenshot shows the 'Spanning Tree > MSTP' configuration page. At the top, there is a 'Step:' dropdown set to '1. Configure Global' and an 'Action:' dropdown 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	32768.0030F1245660
Bridge ID	20	Root Port	2
Max Age	15 sec	Root Path Cost	32768.000001010010
Hello Time	23 sec	Configuration Changes	500000
Forward Delay	2 sec	Last Topology Change	0 days, 1 hours, 10 minutes, 0 seconds

To add additional VLAN groups to an MSTP instance:

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


Figure 132: Adding a VLAN to an MST Instance

The screenshot shows the 'Spanning Tree > MSTP' configuration page with the 'Action:' dropdown 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 133: 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 220](#) for additional information.)
 - **Discarding** – Port receives STA configuration messages, but does not forward packets.
 - **Learning** – Port has transmitted configuration messages for an interval set by the Forward Delay parameter without receiving contradictory information. Port address table is cleared, and the port begins learning addresses.
 - **Forwarding** – Port forwards packets, and continues learning addresses.
- ◆ **Priority** – Defines the priority used for this port in the Spanning Tree Protocol. If the path cost for all ports on a switch are the same, the port with the highest

priority (i.e., lowest value) will be configured as an active link in the Spanning Tree. This makes a port with higher priority less likely to be blocked if the Spanning Tree Protocol is detecting network loops. Where more than one port is assigned the highest priority, the port with lowest numeric identifier will be enabled. (Default: 128; Range: 0-240, in steps of 16)

- ◆ **Admin MST Path Cost** – This parameter is used by the MSTP to determine the best path between devices. Therefore, lower values should be assigned to ports attached to faster media, and higher values assigned to ports with slower media. (Path cost takes precedence over port priority.) Note that when the Path Cost Method is set to short (page 3-63), the maximum path cost is 65,535.

(Range: 0 for auto-configuration, 1-65535 for the short path cost method, 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 8021w standard exceeds 65,535, the default is set to 65,535.

The recommended range is listed in [Table 11 on page 217](#).

The default path costs are listed in [Table 12 on page 217](#).

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 134: Configuring MSTP Interface Settings

Spanning Tree > MSTP

Step: 2. Configure Interface Action: Configure

MST ID: 0

Interface: ☒ Port ☐ Trunk

Spanning Tree Port List Total: 28

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

To display MSTP parameters for a port or trunk:

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

Figure 135: Displaying MSTP Interface Settings

Spanning Tree > MSTP

Step: 2. Configure Interface Action: Show Information

MST ID: 0

Interface: ☒ Port ☐ Trunk

Spanning Tree Port List Total: 28

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

Congestion Control

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

Congestion Control includes following options:

- ◆ **Rate Limiting** – Sets the input and output rate limits for a port.
- ◆ **Storm Control** – Sets the traffic storm threshold for each interface.
- ◆ **Automatic Traffic Control** – Sets thresholds for broadcast and multicast storms which can be used to trigger configured rate limits or to shut down a port.

Rate Limiting

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

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

Parameters

These parameters are displayed:


- ◆ **Interface** – Displays the switch's ports or trunks.
- ◆ **Type** – Indicates the port type. (1000BASE-T, 1000BASE SFP, 10GBASE SFP+)
- ◆ **Status** – Enables or disables the rate limit. (Default: Disabled)
- ◆ **Rate** – Sets the rate limit level.
(Range: 64-1,000,000 per second for 1000BASE-T Ethernet ports;
64 - 10,000,000 kbits per second for 10G Ethernet ports)

Web Interface

To configure rate limits:

1. Click Traffic, Rate Limit.
2. Enable the rate limit Status for the required ports.
3. Set the rate limit for the individual ports.
4. Click Apply.

Figure 136: Configuring Rate Limits



The screenshot shows the 'Traffic > Rate Limit' configuration page. At the top, there's a breadcrumb 'Traffic > Rate Limit'. Below it, a tabbed interface shows 'Port' selected over 'Trunk'. A 'Port Rate Limit List' is displayed with a 'Total: 28' and three numbered tabs (1, 2, 3). The table has columns for Port, Type, Status, Input Rate (kbits/sec), and Output Rate (kbits/sec). The input and output rate columns have a range of (64-10000000). All five ports shown are 1000BASE-T and have their status set to 'Enabled' with a checkbox.

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

Storm Control

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

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

Command Usage

- ◆ Broadcast Storm Control is enabled by default.
- ◆ When traffic exceeds the threshold specified for broadcast and multicast or unknown unicast traffic, packets exceeding the threshold are dropped until the rate falls back down beneath the threshold.
- ◆ Traffic storms can be controlled at the hardware level using Storm Control or at the software level using [Automatic Traffic Control](#) which triggers various control responses. However, only one of these control types can be applied to a

port. Enabling hardware-level storm control on a port will disable automatic storm control on that port.

- ◆ Rate limits set by the storm control function are also used by automatic storm control when the control response is set to rate control on the Auto Traffic Control (Configure Interface) page.
- ◆ Using both rate limiting and storm control on the same interface may lead to unexpected results. It is therefore not advisable to use both of these features on the same interface.

Parameters

These parameters are displayed:

- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Type** – Indicates interface type. (1000BASE-T, 1000BASE SFP, 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: Enabled for broadcast storm control, disabled for multicast and unknown unicast storm control)
- ◆ **Rate** – Threshold level as a rate; i.e., packets per second. (Range: 500-14880000 pps; Default: Disabled for unknown unicast and multicast traffic, 500 pps for broadcast traffic)

Web Interface

To configure broadcast storm control:

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

Figure 137: Configuring Storm Control

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


Automatic Traffic Control

Use the Traffic > Congestion Control > 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 138: 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 237](#)).
- ◆ 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 139: 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 232](#)). 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 237](#)).

Parameters

These parameters are displayed:

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

Web Interface

To configure the response timers for automatic storm control:

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

Figure 140: Configuring ATC Timers

Traffic > Auto Traffic Control

Step: 1. Configure Global ▼

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

Configuring ATC Thresholds and Responses

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

Parameters

These parameters are displayed:

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

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

- ◆ **Auto Release Control** – Automatically stops a traffic control response of rate limiting when traffic falls below the alarm clear threshold and the release timer expires as illustrated in [Figure 138 on page 234](#). 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 141: Configuring ATC Interface Attributes

Traffic > Auto Traffic Control

Step: 2. Configure Interface

Storm Control ☒ Broadcast ☐ Multicast

Auto Traffic Control Broadcast List Total: 28

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

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

This chapter describes the following basic topics:

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

Layer 2 Queue Settings

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

Setting the Default Priority for Interfaces

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

Command Usage

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

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

Parameters

These parameters are displayed:


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

Web Interface

To configure the queue mode:

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

Figure 142: Setting the Default Port Priority



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

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

Command Usage

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

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

- ◆ If Strict and WRR mode is selected, a combination of strict service is used for the high priority queues and weighted service for the remaining queues. The queues assigned to use strict priority should be specified using the Strict Mode field parameter.

- ◆ A weight can be assigned to each of the weighted queues (and thereby to the corresponding traffic priorities). This weight sets the frequency at which each queue is polled for service, and subsequently affects the response time for software applications assigned a specific priority value.

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

- ◆ The specified queue mode applies to all interfaces.

Parameters

These parameters are displayed:

- ◆ **Port** – List of ports. (Range: 1-28/52)
- ◆ **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 and weighted service is used as specified for each queue. Use this parameter to specify the queues assigned to use strict priority when using the strict-weighted queuing mode. (Default: Disabled)
- ◆ **Weight** – Sets a weight for each queue which is used by the WRR scheduler. (Range: 1-15; Default: Weights 1, 2, 4, 6, 8, 10, 12 and 14 are assigned to queues 0 - 7 respectively)

Web Interface

To configure the queue mode:

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

Figure 143: 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. There are 'Apply' and 'Revert' buttons at the bottom right.

Figure 144: Setting the Queue Mode (WRR)

The screenshot shows the 'Traffic > Priority > Queue' configuration page. The 'Port' is set to '1' and the 'Queue Mode' is set to 'WRR'. Below the settings is a 'Queue Setting Table' with 8 rows. The table has two columns: 'Queue ID' and 'Weight (1-127)'. The weights are set to 1, 2, 4, 6, 8, 10, 12, and 14 for Queue IDs 0 through 7 respectively. There are 'Apply' and 'Revert' buttons at the bottom right.

Queue ID	Weight (1-127)
0	1
1	2
2	4
3	6
4	8
5	10
6	12
7	14

Figure 145: Setting the Queue Mode (Strict and WRR)

Traffic > Priority > Queue

Port: 1

Queue Mode: Strict and WRR

Queue Setting Table Total: 8

Queue ID	Strict Mode	Weight (1-127)
0	Disabled	1
1	Disabled	2
2	Disabled	4
3	Disabled	6
4	Disabled	8
5	Disabled	10
6	Disabled	12
7	Disabled	14

Apply Revert

Mapping CoS Values to Egress Queues

Use the Traffic > Priority > PHB to Queue page to specify the hardware output queues to use based on the internal per-hop behavior value. (For more information on exact manner in which the ingress priority tags are mapped to egress queues for internal processing, see [“Mapping CoS Priorities to Internal DSCP Values” on page 252](#)).

The switch processes Class of Service (CoS) priority tagged traffic by using eight priority queues for each port, with service schedules based on strict priority, Weighted Round-Robin (WRR), or a combination of strict and weighted queuing. Up to eight separate traffic priorities are defined in IEEE 802.1p. Default priority levels are assigned according to recommendations in the IEEE 802.1p standard as shown in [Table 13](#). The following table indicates the default mapping of internal per-hop behavior to the hardware queues. The actual mapping may differ if the CoS priorities to internal DSCP values have been modified ([page 252](#)).

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** – List of ports. (Range: 1-28/52)
- ◆ **PHB** – Per-hop behavior, or the priority used for this router hop. (Range: 0-7, where 7 is the highest priority)
- ◆ **Queue** – Output queue buffer. (Range: 0-7, where 7 is the highest CoS priority queue)

Web Interface

To map internal PHB to hardware queues:

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

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

4. Click Apply.

Figure 146: Mapping CoS Values to Egress Queues

Traffic > Priority > PHB to Queue

Action: Configure

Port: 1

PHB (0-7):

Queue (0-7):

Apply Revert

To show the internal PHB to hardware queue map:

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

Figure 147: Showing CoS Values to Egress Queue Mapping

Traffic > Priority > PHB to Queue

Action: Show

Port: 1

PHB to Queue Mapping List Total: 8

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

Default Revert

Layer 3/4 Priority Settings

Mapping Layer 3/4 Priorities to CoS Values

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

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



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

Setting Priority Processing to 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 241](#)) 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 241](#)) is used for priority processing.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28/52)
- ◆ **Trust Mode**
 - **CoS** – Maps layer 3/4 priorities using Class of Service values. (This is the default setting.)
 - **DSCP** – Maps layer 3/4 priorities using Differentiated Services Code Point values.
 - **IP Precedence** – Maps layer 3/4 priorities using IP Precedence values.

Web Interface

To configure the trust mode:

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

Figure 148: Setting the Trust Mode

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

Mapping Ingress DSCP Values to Internal DSCP Values

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

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

Command Usage

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

- ◆ This map is only used when the priority mapping mode is set to DSCP (see [page 248](#)), and the ingress packet type is IPv4. Any attempt to configure the DSCP mutation map will not be accepted by the switch, unless the trust mode has been set to DSCP.
- ◆ Two QoS domains can have different DSCP definitions, so the DSCP-to-PHB/Drop Precedence mutation map can be used to modify one set of DSCP values to match the definition of another domain. The mutation map should be applied at the receiving port (ingress mutation) at the boundary of a QoS administrative domain.

Parameters

These parameters are displayed:

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

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

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

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

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

Web Interface

To map DSCP values to internal PHB/drop precedence:

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

Figure 149: Configuring DSCP to DSCP Internal Mapping

Traffic > Priority > DSCP to DSCP

Action: Configure

Port: 1

DSCP (0-63): 1

PHB (0-7): 3

Drop Precedence: 1: Red

Apply Revert

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

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

Figure 150: Showing DSCP to DSCP Internal Mapping

Traffic > Priority > DSCP to DSCP

Action: Show

DSCP to DSCP Mapping List Total: 64

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

Default Revert

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 252](#).
- ◆ 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 for controlling traffic congestion. (Range: 0 - Green, 3 - Yellow, 1 - Red)

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

CoS	CFI	0	1
0		(0,0)	(0,0)
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 151: 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 152: 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
0	1	0	0
1	0	1	0
1	1	1	0
2	0	2	0
2	1	2	0
3	0	3	0
3	1	3	0
4	0	4	0
4	1	4	0

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

Command Usage

- ◆ Enter any per-hop behavior and drop precedence pair within the internal priority map, and then enter the corresponding CoS/CFI pair.
- ◆ If the packet is forwarded with an 8021.Q tag, the priority value in the egress packet is modified based on the default values shown in [Table 18 on page 254](#), or on the values modified by this function.

Parameters

These parameters are displayed in the web interface:

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

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

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

Web Interface

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

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

Figure 153: Configuring DSCP to CoS Egress Mapping

Traffic > Priority > DSCP to CoS

Action: Configure

Port: 1

PHB (0-7): 1

Drop Precedence: 3: Yellow

CoS (0-7): 1

CFI (0-1): 0

Apply Revert

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

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

Figure 154: Showing DSCP to CoS Egress Mapping

Traffic > Priority > DSCP to CoS

Action: Show

Port: 1

DSCP to CoS Mapping List Total: 24

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

Mapping IP Precedence Values to Internal DSCP Values

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

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

Table 19: Mapping IP Precedence

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

Command Usage

- ◆ Enter per-hop behavior and drop precedence for any of the IP Precedence values 0 - 7.
- ◆ If the priority mapping mode is set the IP Precedence and the ingress packet type is IPv4, then the IP Precedence-to-PHB/Drop Precedence mapping table is used to generate priority and drop precedence values for internal processing.

Parameters

These parameters are displayed in the web interface:

- ◆ **Port** – Port identifier.
- ◆ **IP Precedence** – IP Precedence value in ingress packets. (Range: 0-7)
- ◆ **PHB** – Per-hop behavior, or the priority used for this router hop. (Range: 0-7)

- ◆ **Drop Precedence** – Drop precedence used for controlling traffic congestion. (Range: 0 - Green, 3 - Yellow, 1 - Red)

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

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

Web Interface

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

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

Figure 155: Configuring IP Precedence to DSCP Internal Mapping

The screenshot shows the 'Traffic > DiffServ' configuration page. At the top, there's a breadcrumb 'Traffic > DiffServ'. Below it, the 'Step' is '1. Configure Class' and the 'Action' is 'Add Rule'. The 'Class Name' is 'rd-class' and the 'Type' is 'Match Any'. Under the 'Rule' section, there are three radio buttons: 'ACL', 'IP DSCP (0-63)', and 'IP Precedence (0-7)'. The 'IP DSCP (0-63)' radio button is selected. To the right of the radio buttons, there's a dropdown menu with a downward arrow.

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 156: Showing the IP Precedence to DSCP Internal Map

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

Mapping IP Port Priority to Internal DSCP Values

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

Command Usage

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

Parameters

These parameters are displayed in the web interface:

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

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

Web Interface

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

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

Figure 157: Configuring IP Port Number to DSCP Internal Mapping

The screenshot shows a web interface titled "Traffic > Priority > IP Port to DSCP". Below the title bar, there is an "Action:" dropdown menu set to "Configure". The main configuration area contains the following fields:

- Port:** A dropdown menu with the value "1" selected.
- IP Protocol:** A dropdown menu with the value "TCP" selected.
- Destination Port (0-65535):** A text input field containing the value "21".
- PHB (0-7):** A text input field containing the value "1".
- Drop Precedence:** A dropdown menu with the value "0: Green" selected.

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

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

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

Figure 158: Showing IP Port Number to DSCP Internal Mapping

The screenshot shows a web interface for configuring network settings. At the top, a breadcrumb trail reads "Traffic > Priority > IP Port to DSCP". Below this, there is a section for "Action:" with a dropdown menu set to "Show". Underneath, a "Port" dropdown menu is set to "1". The main content area is titled "IP Port to DSCP Mapping List" with a "Total: 1" indicator. It contains a table with the following data:

	IP Protocol	Destination Port	PHB	Drop Precedence
☐	TCP	21	1	0

At the bottom of the table, there are two buttons: "Delete" and "Revert".

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

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

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

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

Overview

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

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

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



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

Note: You should create a class map before creating a policy map. Otherwise, you will not be able to select a class map from the policy rule settings screen (see [page 265](#)).

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

Configuring a Class Map

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

Command Usage

- ◆ The class map is used with a policy map ([page 265](#)) to create a service policy ([page 275](#)) for a specific interface that defines packet classification, service tagging, and bandwidth policing. Note that one or more class maps can be assigned to a policy map.
- ◆ Up to 32 class maps can be configured.

Parameters

These parameters are displayed:

Add

- ◆ **Class Name** – Name of the class map. (Range: 1-32 characters)
- ◆ **Type** – Only one match command is permitted per class map, so the match-any field refers to the criteria specified by the lone match command.
- ◆ **Description** – A brief description of a class map. (Range: 1-64 characters)

Add Rule

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

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

Traffic > DiffServ

Step: 1. Configure Class Action: Add

Class Name: rd-class

Type: Match Any

Description: class for software group

Apply Revert

To show the configured class maps:

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

Figure 160: Showing Class Maps

The screenshot shows a web interface for configuring DiffServ. At the top, there's a breadcrumb 'Traffic > DiffServ' and a green help icon. Below this, there are two dropdown menus: 'Step: 1. Configure Class' and 'Action: Show'. The main section is titled 'Class List' with a 'Total: 1' indicator. It contains a table with the following data:

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

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

To edit the rules for a class map:

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

Figure 161: Adding Rules to a Class Map

The screenshot shows the 'Traffic > DiffServ' configuration page. The 'Step' is '1. Configure Class' and the 'Action' is 'Add Rule'. The 'Class Name' is 'rd-class' and the 'Type' is 'Match Any'. Under the 'Rule' section, several options are listed with radio buttons: ACL, IP DSCP (0-63), IP Precedence (0-7), IPv6 DSCP (0-63), VLAN ID (1-4094), and CoS (0-7). Each option has a corresponding input field. The 'IP DSCP (0-63)' option is selected. At the bottom right, there are 'Apply' and 'Revert' buttons.

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

The screenshot shows the 'Traffic > DiffServ' configuration page. The 'Step' is '1. Configure Class' and the 'Action' is 'Show Rule'. The 'Class Name' is 'rd-class' and the 'Type' is 'Match Any'. Below the 'Rule List' section, there is a table with 2 rules. The table has a 'Rule' column and a 'Delete' column. The rules are 'IP DSCP 3' and 'IP Precedence 3'. At the bottom right, there are 'Delete' and 'Revert' buttons.

Rule	Delete
IP DSCP 3	☐
IP Precedence 3	☐

Creating QoS Policies

Use the Traffic > DiffServ (Configure Policy) page to create a policy map that can be attached to multiple interfaces. A policy map is used to group one or more class map statements ([page 262](#)), modify service tagging, and enforce bandwidth policing. A policy map can then be bound by a service policy to one or more interfaces ([page 275](#)).

Configuring QoS policies requires several steps. A class map must first be configured which indicates how to match the inbound packets according to an access list, a DSCP or IP Precedence value, or a member of specific VLAN. A policy

map is then configured which indicates the boundary parameters used for monitoring inbound traffic, and the action to take for conforming and non-conforming traffic. A policy map may contain one or more classes based on previously defined class maps.

The class of service or per-hop behavior (i.e., the priority used for internal queue processing) can be assigned to matching packets. In addition, the flow rate of inbound traffic can be monitored and the response to conforming and non-conforming traffic based by one of three distinct policing methods as described below.

Police Flow Meter – Defines the committed information rate (maximum throughput), committed burst size (burst rate), and the action to take for conforming and non-conforming traffic.

Policing is based on a token bucket, where bucket depth (that is, the maximum burst before the bucket overflows) is specified by the “burst” field (BC), and the average rate tokens are removed from the bucket is specified by the “rate” option (CIR). Action may be taken for traffic conforming to the maximum throughput, or exceeding the maximum throughput.

srTCM Police Meter – Defines an enforcer for classified traffic based on a single rate three color meter scheme defined in RFC 2697. This metering policy monitors a traffic stream and processes its packets according to the committed information rate (CIR, or maximum throughput), committed burst size (BC, or burst rate), and excess burst size (BE). Action may taken for traffic conforming to the maximum throughput, exceeding the maximum throughput, or exceeding the excess burst size.

- ◆ The PHB label is composed of five bits, three bits for per-hop behavior, and two bits for the color scheme used to control queue congestion. In addition to the actions defined by this command to transmit, remark the DSCP service value, or drop a packet, the switch will also mark the two color bits used to set the drop precedence of a packet. A packet is marked green if it doesn't exceed the committed information rate and committed burst size, yellow if it does exceed the committed information rate and committed burst size, but not the excess burst size, and red otherwise.
- ◆ The meter operates in one of two modes. In the color-blind mode, the meter assumes that the packet stream is uncolored. In color-aware mode the meter assumes that some preceding entity has pre-colored the incoming packet stream so that each packet is either green, yellow, or red. The marker (re)colors an IP packet according to the results of the meter. The color is coded in the DS field [RFC 2474] of the packet.
- ◆ The behavior of the meter is specified in terms of its mode and two token buckets, C and E, which both share the common rate CIR. The maximum size of the token bucket C is BC and the maximum size of the token bucket E is BE.

The token buckets C and E are initially full, that is, the token count $T_c(0) = BC$ and the token count $T_e(0) = BE$. Thereafter, the token counts T_c and T_e are updated CIR times per second as follows:

- If T_c is less than BC , T_c is incremented by one, else
- if T_e is less than BE , T_e is incremented by one, else
- neither T_c nor T_e is incremented.

When a packet of size B bytes arrives at time t , the following happens if srTCM is configured to operate in Color-Blind mode:

- If $T_c(t) - B \geq 0$, the packet is green and T_c is decremented by B down to the minimum value of 0, else
- if $T_e(t) - B \geq 0$, the packet 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 packet 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 meter (re)colors an IP packet according to the results of the meter. The color is coded in the DS field [RFC 2474] of the packet.
- ◆ The behavior of the meter is specified in terms of its mode and two token buckets, P and C, which are based on the rates PIR and CIR, respectively. The maximum size of the token bucket P is BP and the maximum size of the token bucket C is BC.

The token buckets P and C are initially (at time 0) full, that is, the token count $Tp(0) = BP$ and the token count $Tc(0) = BC$. Thereafter, the token count Tp is incremented by one PIR times per second up to BP and the token count Tc is incremented by one CIR times per second up to BC.

When a packet of size B bytes arrives at time t, the following happens if trTCM is configured to operate in Color-Blind mode:

- If $Tp(t) - B < 0$, the packet is red, else
- if $Tc(t) - B < 0$, the packet is yellow and Tp is decremented by B, else
- the packet is green and both Tp and Tc are decremented by B.

When a packet of size B bytes arrives at time t, the following happens if trTCM is configured to operate in Color-Aware mode:

- If the packet has been precolored as red or if $Tp(t) - B < 0$, the packet is red, else
 - if the packet has been precolored as yellow or if $Tc(t) - B < 0$, the packet is yellow and Tp is decremented by B, else
 - the packet is green and both Tp and Tc are decremented by B.
- ◆ The trTCM can be used to mark a IP packet stream in a service, where different, decreasing levels of assurances (either absolute or relative) are given to packets which are green, yellow, or red. Refer to RFC 2698 for more information on other aspects of trTCM.

Command Usage

- ◆ A policy map can contain 16 class statements that can be applied to the same interface ([page 275](#)). 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 275](#)) 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: 1-256 characters)

Add Rule

- ◆ **Policy Name** – Name of policy map.
- ◆ **Class Name** – Name of a class map that defines a traffic classification upon which a policy can act. A policy map can contain up to 16 class maps.
- ◆ **Action** – This attribute is used to set an internal QoS value in hardware for matching packets. The PHB label is composed of five bits, three bits for per-hop behavior, and two bits for the color scheme used to control queue congestion with the srTCM and trTCM metering functions.
 - **Set CoS** – Configures the service provided to ingress traffic by setting an internal CoS value for a matching packet (as specified in rule settings for a class map). (Range: 0-7)
See [Table 17, “Default Mapping of CoS/CFI to Internal PHB/Drop Precedence,” on page 252](#).
 - **Set PHB** – Configures the service provided to ingress traffic by setting the internal per-hop behavior for a matching packet (as specified in rule settings for a class map). (Range: 0-7)
See [Table 17, “Default Mapping of CoS/CFI to Internal PHB/Drop Precedence,” on page 252](#).
 - **Set IP DSCP** – Configures the service provided to ingress traffic by setting an IP DSCP value for a matching packet (as specified in rule settings for a class map). (Range: 0-63)
- ◆ **Meter** – Check this to define the maximum throughput, burst rate, and the action that results from a policy violation.
- ◆ **Meter Mode** – Selects one of the following policing methods.
 - **Flow** (Police Flow) – Defines the committed information rate (CIR, or maximum throughput), committed burst size (BC, or burst rate), and the action to take for conforming and non-conforming traffic. Policing is based on a token bucket, where bucket depth (that is, the maximum burst before the bucket overflows) is specified by the “burst” field, and the average rate tokens are removed from the bucket is by specified by the “rate” option.

- **Committed Information Rate (CIR)** – Committed rate in kilobits per second. (Range: 0-10000000 kbps or maximum port speed, whichever is lower)
The rate cannot exceed the configured interface speed.
- **Committed Burst Size (BC)** – Committed burst in bytes. (Range: 4000-16000000 bytes)
The burst size cannot exceed 16 Mbytes.
- **Conform** – Specifies that traffic conforming to the maximum committed rate (CIR) and committed burst size (BC) will be transmitted without any change to the DSCP service level.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63)
 - **Transmit** – Transmits in-conformance traffic without any change to the DSCP service level.
- **Violate** – Specifies whether the traffic that exceeds the maximum rate (CIR) will be dropped or the DSCP service level will be reduced.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63)
 - **Drop** – Drops out of conformance traffic.
- **srTCM (Police Meter)** – Defines the committed information rate (CIR, or maximum throughput), committed burst size (BC, or burst rate) and excess burst size (BE), and the action to take for traffic conforming to the maximum throughput, exceeding the maximum throughput but within the excess burst size, or exceeding the excess burst size. In addition to the actions defined by this command to transmit, remark the DSCP service value, or drop a packet, the switch will also mark the two color bits used to set the drop precedence of a packet.

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

- **Committed Information Rate (CIR)** – Committed rate in kilobits per second. (Range: 0-10000000 kbps or maximum port speed, whichever is lower)
The rate cannot exceed the configured interface speed.
- **Committed Burst Size (BC)** – Committed burst in bytes. (Range: 4000-16000000 byte)
The burst size cannot exceed 16 Mbytes.

- **Excess Burst Size (BE)** – Burst in excess of committed burst size.
(Range: 4000-16000000 bytes)

The burst size cannot exceed 16 Mbytes.
- **Conform** – Specifies that traffic conforming to the maximum committed rate (CIR) and committed burst size (BC) will be transmitted without any change to the DSCP service level.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63)
 - **Transmit** – Transmits in-conformance traffic without any change to the DSCP service level.
- **Exceed** – Specifies whether traffic that exceeds the committed maximum rate (CIR) but is within the excess burst size (BE) will be dropped or the DSCP service level will be reduced.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63)
 - **Drop** – Drops out of conformance traffic.
- **Violate** – Specifies whether the traffic that exceeds the excess burst size (BE) will be dropped or the DSCP service level will be reduced.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63)
 - **Drop** – Drops out of conformance traffic.
- **trTCM (Police Meter)** – Defines the committed information rate (CIR, or maximum throughput), peak information rate (PIR), and their associated burst sizes – committed burst size (BC, or burst rate) and peak burst size (BP), and the action to take for traffic conforming to the maximum throughput, exceeding the maximum throughput but within the peak information rate, or exceeding the peak information rate. In addition to the actions defined by this command to transmit, remark the DSCP service value, or drop a packet, the switch will also mark the two color bits used to set the drop precedence of a packet.

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

- **Committed Information Rate (CIR)** – Committed rate in kilobits per second. (Range: 0-10000000 kilobits or maximum port speed, whichever is lower)

The rate cannot exceed the configured interface speed.
- **Committed Burst Size (BC)** – Committed burst in bytes.
(Range: 4000-16000000 bytes)

The burst size cannot exceed 16 Mbytes.

- **Peak Information Rate (PIR)** – Rate in kilobits per second.
(Range: 0-10000000 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: 4000-128000000 bytes)
The rate cannot exceed the configured interface speed.
- **Conform** – Specifies that traffic conforming to the committed maximum rate (CIR) and peak burst size (BC) will be transmitted without any change to the DSCP service level.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63).
 - **Transmit** – Transmits in-conformance traffic without any change to the DSCP service level.
- **Exceed** – Specifies whether traffic that exceeds the committed maximum rate (CIR) or committed burst size (BC) but is within the peak information rate (PIR) will be dropped or the DSCP service level will be reduced.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63).
 - **Drop** – Drops out of conformance traffic.
- **Violate** – Specifies whether the traffic that exceeds the peak information rate (PIR) will be dropped or the DSCP service level will be reduced.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63).
 - **Drop** – Drops out of conformance traffic.

Web Interface

To configure a policy map:

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

6. Click Add.

Figure 163: Configuring a Policy Map

Traffic > DiffServ

Step: 2. Configure Policy Action: Add

Policy Name: rd-policy

Description: for the software group

Apply Revert

To show the configured policy maps:

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

Figure 164: Showing Policy Maps

Traffic > DiffServ

Step: 2. Configure Policy Action: Show

Policy List Total: 1

	Policy Name	Description
☐	rd-policy	for the software group

Delete Revert

To edit the rules for a policy map:

1. Click Traffic, DiffServ.
2. Select Configure Policy from the Step list.
3. Select Add Rule from the Action list.
4. Select the name of a policy map.
5. Set the CoS or per-hop behavior for matching packets to specify the quality of service to be assigned to the matching traffic class. Use one of the metering options to define parameters such as the maximum throughput and burst rate. Then specify the action to take for conforming traffic, the action to take for traffic in excess of the maximum rate but within the peak information rate, or the action to take for a policy violation.
6. Click Apply.

Figure 165: Adding Rules to a Policy Map

Traffic > DiffServ

Step: 2. Configure Policy Action: Add Rule

Policy Name: rd-policy

Rule:

Class Name: rd-class

☒ Action: Set CoS (0-7)

☒ Meter

Meter Mode: Flow

Committed Information Rate (0-10000000): 1000000 kbps

Committed Burst Size (4000-16000000): 4000 bytes

Excess Burst Size (4000-16000000): bytes

Peak Information Rate (0-10000000): kbps

Peak Burst Size (4000-16000000): bytes

Conform: Set IP DSCP (0-63) 0

Exceed: Set IP DSCP (0-63)

Violate: Set IP DSCP (0-63) 0

Apply Revert

To show the rules for a policy map:

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

Figure 166: Showing the Rules for a Policy Map

Traffic > DiffServ

Step: 2. Configure Policy Action: Show Rule

Policy Name: rd

Rule List Total: 1

	Class Name	Action	Meter Mode	Committed Information Rate (kbps)	Committed Burst Size (bytes)	Excess Burst Size (bytes)	Peak Information Rate (kbps)	Peak Burst Size (bytes)	Conform	Exceed	Violate
☐	rd	Set CoS 3	Flow	1000000	4000	N/A	N/A	N/A	Transmit		Drop

Delete Revert

Attaching a Policy Map to a Port

Use the Traffic > DiffServ (Configure Interface) page to bind a policy map to a port.

Command Usage

- ◆ First define a class map, define a policy map, and then bind the service policy to the required interface.
- ◆ Only one policy map can be bound to an interface.
- ◆ The switch does not allow a policy map to be bound to an interface for egress traffic.

Parameters

These parameters are displayed:

- ◆ **Port** – Specifies a port.
- ◆ **Ingress** – Applies the selected rule to ingress traffic.
- ◆ **Egress** – Applies the selected rule to egress traffic.

Web Interface

To bind a policy map to a port:

1. Click Traffic, DiffServ.
2. Select Configure Interface from the Step list.
3. Check the box under the Ingress 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

Traffic > DiffServ

Step: 3. Configure Interface

Port Service Policy List Total: 28

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

Apply Revert

VoIP Traffic Configuration

This chapter covers the following topics:

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

Overview

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

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

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

Configuring VoIP Traffic

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

Command Usage

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

Parameters

These parameters are displayed:

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



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

Web Interface

To configure global settings for a Voice VLAN:

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

Figure 168: Configuring a Voice VLAN

Traffic > VoIP

Step: 1. Configure Global

Auto Detection Status ☒ Enabled

Voice VLAN 1234

Voice VLAN Aging Time (5-43200) 3000 sec

Apply Revert

Configuring Telephony OUI

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

Parameters

These parameters are displayed:

- ◆ **Telephony OUI** – Specifies a MAC address range to add to the list.
(Format: xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx)
- ◆ **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 Action: Add

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

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

Description: old phones

Apply Revert

To show the MAC OUI numbers used for VoIP equipment:

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

Figure 170: Showing an OUI Telephony List

Traffic > VoIP

Step: 2. Configure OUI Action: Show

Telephony OUI List Total: 3

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

Delete Revert

Configuring VoIP Traffic Ports

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

Command Usage

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

Parameters

These parameters are displayed:

- ◆ **Mode** – Specifies if the port will be added to the Voice VLAN when VoIP traffic is detected. (Default: None)
 - **None** – The Voice VLAN feature is disabled on the port. The port will not detect VoIP traffic or be added to the Voice VLAN.
 - **Auto** – The port will be added as a tagged member to the Voice VLAN when VoIP traffic is detected on the port. You must select a method for detecting VoIP traffic, either OUI or 802.1AB (LLDP). When OUI is selected, be sure to configure the MAC address ranges in the Telephony OUI list.
 - **Manual** – The Voice VLAN feature is enabled on the port, but the port must be manually added to the Voice VLAN.
- ◆ **Security** – Enables security filtering that discards any non-VoIP packets received on the port that are tagged with the voice VLAN ID. VoIP traffic is identified by source MAC addresses configured in the Telephony OUI list, or through LLDP that discovers VoIP devices attached to the switch. Packets received from non-VoIP sources are dropped. (Default: Disabled)
- ◆ **Discovery Protocol** – Selects a method to use for detecting VoIP traffic on the port. (Default: OUI)
 - **OUI** – Traffic from VoIP devices is detected by the Organizationally Unique Identifier (OUI) of the source MAC address. OUI numbers are assigned to vendors and form the first three octets of a device MAC address. MAC address OUI numbers must be configured in the Telephony OUI list so that the switch recognizes the traffic as being from a VoIP device.
 - **LLDP** – Uses LLDP (IEEE 802.1AB) to discover VoIP devices attached to the port. LLDP checks that the “telephone bit” in the system capability TLV is turned on. See [“Link Layer Discovery Protocol” on page 401](#) 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



Traffic > VoIP

Step: 3. Configure Interface

VoIP Port List Total: 28

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

Security Measures

You can configure this switch to authenticate users logging into the system for management access using local or remote authentication methods. Port-based authentication using IEEE 802.1X can also be configured to control either management access to the uplink ports or client access to the data ports. This switch provides secure network management access using the following options:

- ◆ [AAA](#) – Use local or remote authentication to configure access rights, specify authentication servers, configure remote authentication and accounting.
- ◆ [User Accounts](#) – Manually configure access rights on the switch for specified users.
- ◆ [Web Authentication](#) – Allows stations to authenticate and access the network in situations where 802.1X or Network Access authentication methods are infeasible or impractical.
- ◆ [Network Access](#) – Configure MAC authentication, intrusion response, dynamic VLAN assignment, and dynamic QoS assignment.
- ◆ [HTTPS](#) – Provide a secure web connection.
- ◆ [SSH](#) – Provide a secure shell (for secure Telnet access).
- ◆ [ACL](#) – Access Control Lists provide packet filtering for IP frames (based on address, protocol, Layer 4 protocol port number or TCP control code).
- ◆ [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.
- ◆ [DHCPv4 Snooping](#) – Filter IPv4 traffic on insecure ports for which the source address cannot be identified via DHCPv4 snooping.
- ◆ [DHCPv6 Snooping](#) – Filter IPv6 traffic on insecure ports for which the source address cannot be identified via DHCPv6 snooping.
- ◆ [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 285](#).
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




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: 1-65535; Default: 5)
 - **Authentication Retries** – Number of times the switch tries to authenticate logon access via the authentication server. (Range: 1-30; Default: 2)

- **Set Key** – Mark this box to set or modify the encryption key.
- **Authentication Key** – Encryption key used to authenticate logon access for client. Enclose any string containing blank spaces in double quotes. (Maximum length: 48 characters)
- **Confirm Authentication Key** – Re-type the string entered in the previous field to ensure no errors were made. The switch will not change the encryption key if these two fields do not match.

◆ TACACS+

- **Global** – Provides globally applicable TACACS+ settings.
- **Server Index** – Specifies the index number of the server to be configured. The switch currently supports only one TACACS+ server.
- **Server IP Address** – Address of the TACACS+ server. (A Server Index entry must be selected to display this item.)
- **Authentication Server TCP Port** – Network (TCP) port of TACACS+ server used for authentication messages. (Range: 1-65535; Default: 49)
- **Authentication Timeout** – The number of seconds the switch waits for a reply from the TACACS+ server before it resends the request. (Range: 1-540; Default: 5)
- **Authentication Retries** – Number of times the switch tries to authenticate logon access via the authentication server. (Range: 1-30; Default: 2)
- **Set Key** – Mark this box to set or modify the encryption key.
- **Authentication Key** – Encryption key used to authenticate logon access for client. Enclose any string containing blank spaces in double quotes. (Maximum length: 48 characters)
- **Confirm Authentication Key** – Re-type the string entered in the previous field to ensure no errors were made. The switch will not change the encryption key if these two fields do not match.

Configure Group

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

When specifying the priority sequence for a sever, the server index must already be defined (see [“Configuring Local/Remote Logon Authentication” on page 285](#)).

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)

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

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

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

Figure 175: Configuring Remote Authentication Server (TACACS+)

Security > AAA > Server

Step: 1. Configure Server

Server Type: ☐ RADIUS ☒ TACACS+

☐ Global | Server Index: ☒ 1

Server IP Address: 10.20.30.40

Authentication Server TCP Port (1-65535): 200

Authentication Timeout (1-540): 10 sec

Authentication Retries (1-30): 5

☒ Set Key

Authentication Key:

Confirm Authentication Key:

Apply Revert

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

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

Figure 176: Configuring AAA Server Groups

Security > AAA > Server

Step: 2. Configure Group Action: 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

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

Buttons: 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.
 - **Exec** – Administrative accounting for local console, Telnet, or SSH connections.

- ◆ **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 285](#)). 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** – Accounting for end users.

- **Port** – Port identifier. (Range: 1-28/52)
- **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)

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.

- ◆ **Exec** – Administrative accounting 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.

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.

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 trail at the top is "Security > AAA > Accounting". Below this, there is a "Step:" dropdown menu currently set to "1. Configure Global". The main configuration area has 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, 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: 802.1X

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: 2

	Accounting Type	Method Name	Accounting Notice	Server Group Name
☐	802.1X	default	Start-Stop	radius
☐	EXEC	default	Start-Stop	tacacs+

Delete Revert

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

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

Figure 181: Configuring AAA Accounting Service for 802.1X Service

The screenshot shows the 'Security > AAA > Accounting' configuration page. The 'Step' dropdown is set to '3. Configure Service'. Under 'Accounting Type', the '802.1X' radio button is selected. Below this, the 'Port Method List' is displayed with a total of 28 methods. A table lists the first five methods, all with the name 'default'.

Port	Method Name
1	default
2	default
3	default
4	
5	

Figure 182: Configuring AAA Accounting Service for Exec Service

The screenshot shows the 'Security > AAA > Accounting' configuration page. The 'Step' dropdown is set to '3. Configure Service'. Under 'Accounting Type', the 'EXEC' radio button is selected. Below this, there are input fields for 'Console Method Name' and 'VTY Method Name', both containing the value 'default'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 183: Displaying a Summary of Applied AAA Accounting Methods


The screenshot shows the 'Security > AAA > Accounting' configuration page. The 'Step' dropdown is set to '4. Show Information'. Under 'Summary' and 'Statistics', the 'Summary' radio button is selected. Below this, the 'Method List' is displayed with a total of 2 methods. A table lists the accounting methods and their associated server groups and interfaces.

Accounting Type	Method Name	Server Group Name	Interface
802.1X	default	radius	
EXEC	default	tacacs+	

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

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

Figure 184: Displaying Statistics for AAA Accounting Sessions



Security > AAA > Accounting			
Step: 4. Show Information			
Summary Statistics			
Accounting Statistics Total: 2			
User Name	Accounting Type	Interface	Time Elapsed
Bob	802.1X	Eth1/1	3:44:55
Ted	802.1X	Eth1/5	1:24:51

Configuring AAA Authorization

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

Command Usage

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

Parameters

These parameters are displayed:

Configure Method

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

The group name “tacacs+” specifies all configured TACACS+ hosts (see [“Configuring Local/Remote Logon Authentication” on page 285](#)). 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 185: Configuring AAA Authorization Methods

The screenshot shows the 'Security > AAA > Authorization' configuration page. At the top, there is a breadcrumb trail. Below it, a 'Step:' dropdown is set to '1. Configure Method' and an 'Action:' dropdown is set to 'Add'. The main configuration area contains three fields: 'Authorization Type' with a dropdown menu showing 'EXEC', 'Method Name' with a text input field containing 'default', and 'Server Group Name' with a radio button selected next to a dropdown menu showing 'tacacs+'. There is also an empty radio button next to an empty text input field. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 186: Showing AAA Authorization Methods

The screenshot shows the 'Security > AAA > Authorization' configuration page. The 'Step' dropdown is set to '1. Configure Method' and the 'Action' dropdown is set to 'Show'. Below this, the 'Method List' is displayed with a total of 3 items. The list is a table with columns for checkboxes, Authorization Type, Method Name, and Server Group Name.

	Authorization Type	Method Name	Server Group Name
☐	EXEC	default	tacacs+
☐	EXEC	console	radius
☐	EXEC	telnet	tacacs+

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

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 187: Configuring AAA Authorization Methods for Exec Service

The screenshot shows the 'Security > AAA > Authorization' configuration page with the 'Step' dropdown set to '2. Configure Service'. The 'Authorization Type' is set to 'EXEC'. The 'Console Method Name' and 'VTY Method Name' fields both contain 'tps-auth'. 'Apply' and 'Revert' buttons are at the bottom right.

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

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

Figure 188: Displaying the Applied AAA Authorization Method



The screenshot shows a web interface for configuring AAA authorization. The breadcrumb trail is "Security > AAA > Authorization". Below the breadcrumb, there is a "Step:" dropdown menu set to "3. Show Information". The main content area is titled "Method List" with a "Total: 3" indicator. It contains a table with four columns: "Authorization Type", "Method Name", "Server Group Name", and "Interface".

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

Configuring User Accounts

Use the Security > User Accounts page to control management access to the switch based on manually configured user names and passwords.

Command Usage

- ◆ The default guest name is "guest" with the password "guest." The default administrator name is "admin" with the password "admin."
- ◆ The guest only has read access for most configuration parameters. However, the administrator has write access for all parameters governing the onboard agent. You should therefore assign a new administrator password as soon as possible, and store it in a safe place.

Parameters

These parameters are displayed:

- ◆ **User Name** – The name of the user.
(Maximum length: 32 characters; maximum number of users: 16)
- ◆ **Access Level** – Specifies command access privileges. (Range: 0-15)
Level 0, 8 and 15 are designed for users (guest), managers (network maintenance), and administrators (top-level access). The other levels can be used to configured specialized access profiles.
Level 0-7 provide the same default access to a limited number of commands which display the current status of the switch, as well as several database clear and reset functions. These commands are equivalent to those available under Normal Exec command mode in the CLI.
Level 8-14 provide the same default access privileges, including additional commands beyond those provided for Levels 0-7 (equivalent to CLI Normal Exec command mode), and a subset of the configuration commands provided for Level 15 (equivalent to CLI Privileged Exec command mode).
Level 15 provides full access to all commands.
The privilege level associated with any command can be changed using the "privilege" command described in the *CLI Reference Guide*.

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

◆ **Password Type** – Specifies the following options:

- **No Password** – No password is required for this user to log in.
- **Plain Password** – Plain text unencrypted password.
- **Encrypted Password** – Encrypted password.

The encrypted password is required for compatibility with legacy password settings (i.e., plain text or encrypted) when reading the configuration file during system bootup or when downloading the configuration file from a TFTP or FTP server. There is no need for you to manually configure encrypted passwords.

◆ **Password** – Specifies the user password. (Range: 0-32 characters, case sensitive)

◆ **Confirm Password** – Re-type the string entered in the previous field to ensure no errors were made. The switch will not change the password if these two fields do not match.

Web Interface

To configure user accounts:

1. Click Security, User Accounts.
2. Select Add from the Action list.
3. Specify a user name, select the user's access level, then enter a password if required and confirm it.
4. Click Apply.

Figure 189: Configuring User Accounts

The screenshot shows a web interface titled "Security > User Accounts". Below the title bar, there is an "Action:" dropdown menu with "Add" selected. The main form contains the following fields:


- User Name:** A text input field containing "bob".
- Access Level:** A dropdown menu with "15 (Privileged)" selected.
- Password Type:** A dropdown menu with "Plain Password" selected.
- Password:** A text input field with masked characters (dots).
- Confirm Password:** A text input field with masked characters (dots).

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 190: Showing User Accounts



Security > User Accounts		
Action: Show		
User Account List Total: 3		
	User Name	Access Level
☐	admin	15
☐	guest	0
☐	bob	15

Delete Revert

Web Authentication

Web authentication allows stations to authenticate and access the network in situations where 802.1X or Network Access authentication are infeasible or impractical. The web authentication feature allows unauthenticated hosts to request and receive a DHCP assigned IP address and perform DNS queries. All other traffic, except for HTTP protocol traffic, is blocked. The switch intercepts HTTP protocol traffic and redirects it to a switch-generated web page that facilitates user name and password authentication via RADIUS. Once authentication is successful, the web browser is forwarded on to the originally requested web page. Successful authentication is valid for all hosts connected to the port.



Note: RADIUS authentication must be activated and configured properly for the web authentication feature to work properly. (See [“Configuring Local/Remote Logon Authentication”](#) on page 285.)

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 191: 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 192: Configuring Interface Settings for Web Authentication

Security > Web Authentication

Step: 2. Configure Interface

Port: 1

Status: ☒ Enabled

Apply Revert

Authenticated Host List Total: 2

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

Re-authenticate Revert

Network Access (MAC Address Authentication)

Some devices connected to switch ports may not be able to support 802.1X authentication due to hardware or software limitations. This is often true for devices such as network printers, IP phones, and some wireless access points. The switch enables network access from these devices to be controlled by authenticating device MAC addresses with a central RADIUS server.



Note: RADIUS authentication must be activated and configured properly for the MAC Address authentication feature to work properly. (See [“Configuring Remote Logon Authentication Servers” on page 286.](#))

Note: MAC authentication cannot be configured on trunk ports.

Command Usage

- ◆ MAC address authentication controls access to the network by authenticating the MAC address of each host that attempts to connect to a switch port. Traffic received from a specific MAC address is forwarded by the switch only if the source MAC address is successfully authenticated by a central RADIUS server. While authentication for a MAC address is in progress, all traffic is blocked until authentication is completed. On successful authentication, the RADIUS server may optionally assign VLAN and quality of service settings for the switch port.
- ◆ When enabled on a port, the authentication process sends a Password Authentication Protocol (PAP) request to a configured RADIUS server. The user name and password are both equal to the MAC address being authenticated. On the RADIUS server, PAP user name and passwords must be configured in the MAC address format XX-XX-XX-XX-XX-XX (all in upper case).
- ◆ Authenticated MAC addresses are stored as dynamic entries in the switch secure MAC address table and are removed when the aging time expires. The maximum number of secure MAC addresses supported for the switch system is 1024.
- ◆ Configured static MAC addresses are added to the secure address table when seen on a switch port. Static addresses are treated as authenticated without sending a request to a RADIUS server.
- ◆ When port status changes to down, all MAC addresses mapped to that port are cleared from the secure MAC address table. Static VLAN assignments are not restored.
- ◆ The RADIUS server may optionally return a VLAN identifier list to be applied to the switch port. The following attributes need to be configured on the RADIUS server.
 - **Tunnel-Type** = VLAN
 - **Tunnel-Medium-Type** = 802

▪ **Tunnel-Private-Group-ID** = 1u,2t [VLAN ID list]

The VLAN identifier list is carried in the RADIUS “Tunnel-Private-Group-ID” attribute. The VLAN list can contain multiple VLAN identifiers in the format “1u,2t,3u” where “u” indicates an untagged VLAN and “t” a tagged VLAN.

- ◆ The RADIUS server may optionally return dynamic QoS assignments to be applied to a switch port for an authenticated user. The “Filter-ID” attribute (attribute 11) can be configured on the RADIUS server to pass the following QoS information:

Table 21: Dynamic QoS Profiles

Profile	Attribute Syntax	Example
DiffServ	service-policy-in = <i>policy-map-name</i>	service-policy-in=p1
Rate Limit	rate-limit-input = <i>rate</i>	rate-limit-input=100 (in units of 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 349](#)).

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 193: Configuring Global Settings for Network Access

Security > Network Access

Step: 1. Configure Global

Aging Status ☒ Enabled

Reauthentication Time (120-1000000) sec

Apply Revert

Configuring Network Access for Ports

Use the Security > Network Access (Configure Interface - General) page to configure MAC authentication on switch ports, including enabling address authentication, setting the maximum MAC count, and enabling dynamic VLAN or dynamic QoS assignments.

Parameters

These parameters are displayed:

◆ MAC Authentication

- **Status** – Enables MAC authentication on a port. (Default: Disabled)
- **Intrusion** – Sets the port response to a host MAC authentication failure to either block access to the port or to pass traffic through. (Options: Block, Pass; Default: Block)
- **Max MAC Count**¹³ – 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**¹³ – Sets the maximum number of MAC addresses that can be authenticated on a port interface via all forms of authentication (including Network Access and IEEE 802.1X). (Range: 1-1024; Default: 1024)

13. 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.

- ◆ **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 163](#)). 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 349](#)).

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

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

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

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

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

- ◆ **Dynamic QoS** – Enables dynamic QoS assignment for an authenticated port. (Default: Disabled)
- ◆ **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 194: Configuring Interface Settings for Network Access

Security > Network Access

Step: 2. Configure Interface

General Link Detection

Port List Total: 52

Port	MAC Authentication			Network Access Max MAC Count (1-1024)	Guest VLAN (0-4094, 0: Disabled)	Dynamic VLAN	Dynamic QoS	MAC Filter ID (1-64)
	Status	Intrusion	Max MAC Count (1-1024)					
1	☐ Enabled	Block	1024	1024	0	☒ Enabled	☐ Enabled	☐
2	☒ Enabled	Pass	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 195: Configuring Link Detection for Network Access

Port	Link Detection Status	Condition	Action
1	☐ Enabled	Link down	Trap
2	☒ Enabled	Link up and down	Trap
3	☐ Enabled	Link down	Trap
4	☐ Enabled	Link down	Trap
5	☐ Enabled	Link down	Trap

Configuring a MAC Address Filter

Use the Security > Network Access (Configure MAC Filter) page to designate specific MAC addresses or MAC address ranges as exempt from authentication. MAC addresses present in MAC Filter tables activated on a port are treated as pre-authenticated on that port.

Command Usage

- ◆ Specified MAC addresses are exempt from authentication.
- ◆ Up to 65 filter tables can be defined.
- ◆ There is no limitation on the number of entries used in a filter table.

Parameters

These parameters are displayed:

- ◆ **Filter ID** – Adds a filter rule for the specified filter. (Range: 1-64)
- ◆ **MAC Address** – The filter rule will check ingress packets against the entered MAC address or range of MAC addresses (as defined by the MAC Address Mask).
- ◆ **MAC Address Mask** – The filter rule will check for the range of MAC addresses defined by the MAC bit mask. If you omit the mask, the system will assign the default mask of an exact match. (Range: 000000000000 - FFFFFFFF; Default: FFFFFFFF)

Web Interface

To add a MAC address filter for MAC authentication:

1. Click Security, Network Access.
2. Select Configure MAC Filter from the Step list.
3. Select Add from the Action list.
4. Enter a filter ID, MAC address, and optional mask.
5. Click Apply.

Figure 196: Configuring a MAC Address Filter for Network Access

Security > Network Access

Step: 3. Configure MAC Filter Action: Add

Filter ID (1-64) 22

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

MAC Address Mask FFFFFFFFFF

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 197: 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 198: Showing Addresses Authenticated for Network Access

MAC Address	Interface	RADIUS Server	Time	Attribute
00-00-86-45-F2-23	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 16m 12s	Dynamic
00-00-E8-5E-E1-DD	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 32m 24s	Dynamic
00-00-E8-81-93-30	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 40m 32s	Dynamic
00-01-80-31-B8-30	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 18m 51s	Dynamic
00-01-80-36-95-D8	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 32m 22s	Dynamic

Configuring HTTPS

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

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

Command Usage

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

- ◆ The client and server establish a secure encrypted connection.
A padlock icon should appear in the status bar for Internet Explorer 11, Mozilla Firefox 39, or Google Chrome 44, or more recent versions.

- ◆ The following web browsers and operating systems currently support HTTPS:

Table 22: HTTPS System Support

Web Browser	Operating System
Internet Explorer 11.x or later	Windows 98, Windows NT (with service pack 6a), Windows 2000, XP, Vista, 7, 8, 10
Mozilla Firefox 39 or later	Windows 2000, XP, Vista, 7, 8, 10, Linux
Google Chrome 44 or later	Windows XP, Vista, 7, 8, 10

- ◆ To specify a secure-site certificate, see [“Replacing the Default Secure-site Certificate” on page 315](#).



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

Parameters

These parameters are displayed:

- ◆ **HTTPS Status** – Allows you to enable/disable the HTTPS server feature on the switch. (Default: Enabled)
- ◆ **HTTPS Port** – Specifies the TCP port number used for HTTPS connection to the switch’s web interface. (Range: 1-65535, except for the following reserved ports: 1 and 25 - Linux kernel, 23 - Telnet, 80 - HTTP; Default: Port 443)

Web Interface

To configure HTTPS:

1. Click Security, HTTPS.
2. Select Configure Global from the Step list.
3. Enable HTTPS and specify the port number if required.
4. Click Apply.

Figure 199: Configuring HTTPS

The screenshot shows a web interface for configuring HTTPS. At the top, it says 'Security > HTTPS'. Below that, there's a dropdown menu for 'Action' set to 'Configure Global'. Underneath, 'HTTPS Status' is set to 'Enabled' with a checked checkbox. 'HTTPS Port (1-65535)' is set to '443' in a text box. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 101](#) or type “reload” at the command prompt: `Console#reload`

Parameters

These parameters are displayed:

- ◆ **TFTP Server IP Address** – IP address of TFTP server which contains the certificate file.
- ◆ **Certificate Source File Name** – Name of certificate file stored on the TFTP server.
- ◆ **Private Key Source File Name** – Name of private key file stored on the TFTP server.

- ◆ **Private Password** – Password stored in the private key file. This password is used to verify authorization for certificate use, and is verified when downloading the certificate to the switch.
- ◆ **Confirm Password** – Re-type the string entered in the previous field to ensure no errors were made. The switch will not download the certificate if these two fields do not match.

Web Interface

To replace the default secure-site certificate:

1. Click Security, HTTPS.
2. Select Copy Certificate from the Step list.
3. Fill in the TFTP server, certificate and private key file name, and private password.
4. Click Apply.

Figure 200: Downloading the Secure-Site Certificate

The screenshot shows a web interface for configuring HTTPS. The breadcrumb is "Security > HTTPS". The "Action:" dropdown is set to "Copy Certificate". Below this are five input fields: "TFTP Server IP Address" (192.168.0.4), "Certificate Source File Name" (postal-site-certificate), "Private Key Source File Name" (postal-private-key), "Private Password" (masked with dots), and "Confirm Password" (masked with dots). At the bottom right are "Apply" and "Revert" buttons.

Configuring the Secure Shell

The Berkeley-standard includes remote access tools originally designed for Unix systems. Some of these tools have also been implemented for Microsoft Windows and other environments. These tools, including commands such as *rlogin* (remote login), *rsh* (remote shell), and *rcp* (remote copy), are not secure from hostile attacks.

Secure Shell (SSH) includes server/client applications intended as a secure replacement for the older Berkeley remote access tools. SSH can also provide remote management access to this switch as a secure replacement for Telnet. When the client contacts the switch via the SSH protocol, the switch generates a public-key that the client uses along with a local user name and password for access authentication. SSH also encrypts all data transfers passing between the switch and

SSH-enabled management station clients, and ensures that data traveling over the network arrives unaltered.



Note: You need to install an SSH client on the management station to access the switch for management via the SSH protocol.

Note: The switch supports both SSH Version 1.5 and 2.0 clients.

Command Usage

The SSH server on this switch supports both password and public key authentication. If password authentication is specified by the SSH client, then the password can be authenticated either locally or via a RADIUS or TACACS+ remote authentication server, as specified on the System Authentication page ([page 285](#)). 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 322](#), or use the "copy tftp public-key" command (see the "copy" command in the *CLI Reference Guide*) to copy a file containing the public key for all the SSH client's granted management access to the switch. (Note that these clients must be configured locally on the switch via the User Accounts page as described on [page 299](#).) The clients are subsequently authenticated using these keys. The current firmware only accepts public key files based on standard UNIX format as shown in the following example for an RSA Version 1 key:

```
1024 35
134108168560989392104094492015542534763164192187295892114317388005553616163105
177594083868631109291232226828519254374603100937187721199696317813662774141689
851320491172048303392543241016379975923714490119380060902539484084827178194372
```

```
288402533115952134861022902978982721353267131629432532818915045306393916643  
steve@192.168.1.19
```

4. *Set the Optional Parameters* – On the SSH Settings page, configure the optional parameters, including the authentication timeout, the number of retries, and the server key size.
5. *Enable SSH Service* – On the SSH Settings page, enable the SSH server on the switch.
6. *Authentication* – One of the following authentication methods is employed:
Password Authentication (for SSH v1.5 or V2 Clients)
 - a. The client sends its password to the server.
 - b. The switch compares the client's password to those stored in memory.
 - c. If a match is found, the connection is allowed.



Note: To use SSH with only password authentication, the host public key must still be given to the client, either during initial connection or manually entered into the known host file. However, you do not need to configure the client's keys.

Public Key Authentication – When an SSH client attempts to contact the switch, the SSH server uses the host key pair to negotiate a session key and encryption method. Only clients that have a private key corresponding to the public keys stored on the switch can access it. The following exchanges take place during this process:

Authenticating SSH v1.5 Clients

- a. The client sends its RSA public key to the switch.
- b. The switch compares the client's public key to those stored in memory.
- c. If a match is found, the switch uses its secret key to generate a random 256-bit string as a challenge, encrypts this string with the user's public key, and sends it to the client.
- d. The client uses its private key to decrypt the challenge string, computes the MD5 checksum, and sends the checksum back to the switch.
- e. The switch compares the checksum sent from the client against that computed for the original string it sent. If the two checksums match, this means that the client's private key corresponds to an authorized public key, and the client is authenticated.

Authenticating SSH v2 Clients

- a. The client first queries the switch to determine if DSA public key authentication using a preferred algorithm is acceptable.
- b. If the specified algorithm is supported by the switch, it notifies the client to proceed with the authentication process. Otherwise, it rejects the request.

- c. The client sends a signature generated using the private key to the switch.
- d. When the server receives this message, it checks whether the supplied key is acceptable for authentication, and if so, it then checks whether the signature is correct. If both checks succeed, the client is authenticated.



Note: The SSH server supports up to four 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: A host key pair must be configured on the switch before you can enable the SSH server. See [“Generating the Host Key Pair” on page 320](#).

Parameters

These parameters are displayed:

- ◆ **SSH Server Status** – Allows you to enable/disable the SSH server on the switch. (Default: Disabled)
- ◆ **Version** – The Secure Shell version number. Version 2.0 is displayed, but the switch supports management access via either SSH Version 1.5 or 2.0 clients.
- ◆ **Authentication Timeout** – Specifies the time interval in seconds that the SSH server waits for a response from a client during an authentication attempt. (Range: 1-120 seconds; Default: 120 seconds)
- ◆ **Authentication Retries** – Specifies the number of authentication attempts that a client is allowed before authentication fails and the client has to restart the authentication process. (Range: 1-5 times; Default: 3)
- ◆ **Server-Key Size** – Specifies the SSH server key size. (Range: 512-896 bits; Default: 768)
 - The server key is a private key that is never shared outside the switch.
 - The host key is shared with the SSH client, and is fixed at 1024 bits.

Web Interface

To configure the SSH server:

1. Click Security, SSH.
2. Select Configure Global from the Step list.
3. Enable the SSH server.
4. Adjust the authentication parameters as required.
5. Click Apply.

Figure 201: 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
Server-Key Size (512-896)	768

Apply Revert

Generating the Host Key Pair

Use the Security > SSH (Configure Host Key - Generate) page to generate a host public/private key pair used to provide secure communications between an SSH client and the switch. After generating this key pair, you must provide the host public key to SSH clients and import the client's public key to the switch as described in the section ["Importing User Public Keys" on page 322](#).



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 319](#).

Parameters

These parameters are displayed:

- ◆ **Host-Key Type** – The key type used to generate the host key pair (i.e., public and private keys). (Range: RSA (Version 1), DSA (Version 2), Both; Default: Both)
The SSH server uses RSA or DSA for key exchange when the client first establishes a connection with the switch, and then negotiates with the client to select either DES (56-bit) or 3DES (168-bit) for data encryption.



Note: The switch uses only RSA Version 1 for SSHv1.5 clients and DSA Version 2 for SSHv2 clients.

Web Interface

To generate the SSH host key pair:

1. Click Security, SSH.
2. Select Configure Host Key from the Step list.
3. Select Generate from the Action list.
4. Select the host-key type from the drop-down box.
5. Click Apply.

Figure 202: Generating the SSH Host Key Pair

A screenshot of a web interface for configuring SSH. The breadcrumb trail at the top reads "Security > SSH". Below this, there are two dropdown menus: "Step:" with the value "2. Configure Host Key" and "Action:" with the value "Generate". Underneath these is a "Host-Key Type" label followed by a dropdown menu showing "Both". At the bottom right of the form area are two buttons: "Apply" and "Revert".

Security > SSH	
Step: 2. Configure Host Key ▼	Action: Generate ▼
Host-Key Type	Both ▼
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 host-key type to clear.
5. Click Clear.

Figure 203: Showing the SSH Host Key Pair

Security > SSH

Step: 2. Configure Host Key Action: Show

Public-Key of Host-Key

☒ RSA

```
1024 65537
133939690967068772691132582542091118069946187790543050435991419159456328108424510227003
175806564056474112418102266589753060303978654741958937882157740680884883800146758464583
403760868973720460189697876570213012247127554124148560753905598712710960791931537404447
966097709515788354636277697671184100501439926741
```

☐ DSA

```
ssh-dss
AAAAAB3NzaC1kc3MAAACBAly4Beu1CFOPsrgqEeW89esSV0snL8VSBwmiY7r6n8GksbwJJ6wLA9z+4kUAFKbFG
3WLVHMAldj62o3W777PSPGzgyH1CD4KEWHNHuB0oLuGozeZ5HzpDH76zkYwFzIJuKwS7FntQTD+//1wPCmNyi
cQdwJICTnRs1YYpVh4YKYfAAAAAQDwt5uUQHcC8qaFQ3DhNFPZ9wCX5wAAAIACzKDORA5vv91v6WHYBbc9GX6
PH4V+15Lm+Eg87ji0UAYxUGat1q8071h16RRR5XH44CIOCHC4VJjhS1873weYO/IULCZSAEa9GczPpWhZ9EA1
```

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 299](#)).
- ◆ **User Key Type** – The type of public key to upload.
 - RSA: The switch accepts a RSA version 1 encrypted public key.
 - DSA: The switch accepts a DSA version 2 encrypted public key.

The SSH server uses RSA or DSA for key exchange when the client first establishes a connection with the switch, and then negotiates with the client to select either DES (56-bit) or 3DES (168-bit) for data encryption.

The switch uses only RSA Version 1 for SSHv1.5 clients and DSA Version 2 for SSHv2 clients.

- ◆ **TFTP Server IP Address** – The IP address of the TFTP server that contains the public key file you wish to import.
- ◆ **Source File Name** – The public key file to upload.

Web Interface

To copy the SSH user's public key:

1. Click Security, SSH.
2. Select Configure User Key from the Step list.
3. Select Copy from the Action list.
4. Select the user name and the public-key type from the respective drop-down boxes, input the TFTP server IP address and the public key source file name.
5. Click Apply.

Figure 204: Copying the SSH User's Public Key

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

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

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

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

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

Figure 205: Showing the SSH User's Public Key

The screenshot shows a web-based configuration interface for SSH. At the top, it says "Security > SSH". Below that, there's a "Step:" dropdown set to "3. Configure User Key" and an "Action:" dropdown set to "Show". The "User Name" is set to "admin". Under the heading "Public-Key of User-Key", there are two sections: "RSA" and "DSA". The "RSA" section is selected, showing a long alphanumeric string. The "DSA" section is also visible, showing another long alphanumeric string. A "Clear" button is at the bottom right.

```
1024 65537
1480220937772193109967882191897869076673078824151724992407025121828690162016474445339
1667942450570093339422932545152202043583189918588264701199521464739810320305865029839
8473670516025543210438660758858919850241166441730463579799255182511108846314033955899
732895735202234665421721768844978831196756189309582533

ssh-dss
AAAAB3NzaC1kc3MAAACBAKM4i8EgUe89W+vh1+y4z12YpyK8cpCNz30rhCrivlClKmdgE/fiZPsqrYH/ClAB/
sJlrNAX0sJTUxtb8e2GLk+x8UmQJVuSDdklwZ9ZRoFwWAl0Yyi57V1TK3tUnI9ocCbVJNGckJkXd1uac4OP09
tAPEAuXBuAWGpDAGSmg6pHAAAFQCpzwjn0rSaLiTq53jx1tsj0RILZwAAAIbIL0Cr7GC7ARztGE9dRkT9oh9
uhuiAzcXrAcZTmxhjGsEtMlAtxKm+r1O6pnz2aX9KEzMewJEMuDphOTRnSpMv39XtSW1aSXWtKoGAcQgDsFqK
```

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 source or destination address), 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 256.
- ◆ The maximum number of rules per ACL is 1K.
- ◆ An ACL can have up to 1K rules. However, due to resource restrictions, the average number of rules bound to the ports should not exceed 20.
- ◆ The maximum number of rules that can be bound to the ports is 64 for each of the following list types: MAC ACLs, IP ACLs (including Standard and Extended ACLs), IPv6 Standard ACLs, and IPv6 Extended ACLs.

The maximum number of rules (Access Control Entries, or ACEs) stated above is the worst case scenario. In practice, the switch compresses the ACEs in TCAM (a hardware table used to store ACEs), but the actual maximum number of ACEs

possible depends on too many factors to be precisely determined. It depends on the amount of hardware resources reserved at runtime for this purpose.

Auto ACE Compression is a software feature used to compress all the ACEs of an ACL to utilize hardware resources more efficiently. Without compression, one ACE would occupy a fixed number of entries in TCAM. So if one ACL includes 25 ACEs, the ACL would need $(25 * n)$ entries in TCAM, where “n” is the fixed number of TCAM entries needed for one ACE. When compression is employed, before writing the ACE into TCAM, the software compresses the ACEs to reduce the number of required TCAM entries. For example, one ACL may include 128 ACEs which classify a continuous IP address range like 192.168.1.0~255. If compression is disabled, the ACL would occupy $(128*n)$ entries of TCAM, using up nearly all of the hardware resources. When using compression, the 128 ACEs are compressed into one ACE classifying the IP address as 192.168.1.0/24, which requires only “n” entries in TCAM. The above example is an ideal case for compression. The worst case would be if no any ACE can be compressed, in which case the used number of TCAM entries would be the same as without compression. It would also require more time to process the ACEs.

- ◆ If no matches are found down to the end of the list, the traffic is denied. For this reason, frequently hit entries should be placed at the top of the list. There is an implied deny for traffic that is not explicitly permitted. Also, note that a single-entry ACL with only one deny entry has the effect of denying all traffic. You should therefore use at least one permit statement in an ACL or all traffic will be blocked.

Because the switch stops testing after the first match, the order of the conditions is critical. If no conditions match, the packet will be denied.

The order in which active ACLs are checked is as follows:

1. User-defined rules in IP and MAC ACLs for ingress ports are checked in parallel.
2. Rules within an ACL are checked in the configured order, from top to bottom.
3. If the result of checking an IP ACL is to permit a packet, but the result of a MAC ACL on the same packet is to deny it, the packet will be denied (because the decision to deny a packet has a higher priority for security reasons). A packet will also be denied if the IP ACL denies it and the MAC ACL accepts it.

Showing TCAM Utilization

Use the Security > ACL (Configure ACL - Show TCAM) page to show utilization parameters for TCAM (Ternary Content Addressable Memory), including the number policy control entries in use, the number of free entries, and the overall percentage of TCAM in use.

Command Usage

Policy control entries (PCEs) are used by various system functions which rely on rule-based searches, including Access Control Lists (ACLs), IP Source Guard filter rules, Quality of Service (QoS) processes, QinQ, MAC-based VLANs, VLAN translation, or traps.

For example, when binding an ACL to a port, each rule in an ACL will use two PCEs; and when setting an IP Source Guard filter rule for a port, the system will also use two PCEs.

Parameters

These parameters are displayed:

- ◆ **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 206: Showing TCAM Utilization

Security > ACL

Step: 1. Configure ACL Action: Show TCAM

Pool Capability Code:

AM - MAC ACL, A4 - IPv4 ACL, A6S - IPv6 Standard ACL,
A6E - IPv6 extended ACL, DM - MAC diffServ, D4 - IPv4 diffServ,
D6S - IPv6 standard diffServ, D6E - IPv6 extended diffServ,
AEM - Egress MAC ACL, AE4 - Egress IPv4 ACL,
AE6S - Egress IPv6 standard ACL, AE6E - Egress IPv6 extended ACL,
DEM - Egress MAC diffServ, DE4 - Egress IPv4 diffServ,
DE6S - Egress IPv6 standard diffServ,
DE6E - Egress IPv6 extended diffServ, W - Web authentication,
I - IP source guard, I6 - IPv6 source guard, C - CPU interface,
L - Link local, Reserved - Reserved,
ALL - All supported function,

TCAM List Total: 12

Unit	Device	Pool	Total	Used	Free	Capability
1	0	0	128	0	128	A6S D6S
1	0	1	128	0	128	A6E D6E C L
1	0	2	128	0	128	A4 D4
1	0	3	128	0	128	AM DM
1	0	4	128	128	0	I
1	0	5	64	64	0	C
1	0	6	128	128	0	I6
1	0	7	128	128	0	W
1	0	8	128	0	128	AE6S DE6S
1	0	9	128	0	128	AE6E DE6E
1	0	10	128	0	128	AE4 DE4
1	0	11	128	0	128	AEM DEM

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 385](#)).

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 207: Creating an ACL

The screenshot shows a web interface for configuring an ACL. The breadcrumb is "Security > ACL". Below it, there are two dropdown menus: "Step:" set to "2. Configure ACL" and "Action:" set to "Add". The main configuration area has two fields: "ACL Name" with the value "R&D" and "Type" with the value "IP Standard". At the bottom right, there are two buttons: "Apply" and "Revert".

To show a list of ACLs:

1. Click Security, ACL.
2. Select Configure ACL from the Step list.
3. Select Show from the Action list.

Figure 208: Showing a List of ACLs

Security > ACL

Step: 2. Configure ACL Action: Show

ACL List Total: 10

	ACL Name	Type
☐	aciStandard1	IP Standard
☐	aciStandard2	IP Standard
☐	aciExtended1	IP Extended
☐	aciExtended2	IP Extended
☐	aciMAC1	MAC
☐	aciMAC2	MAC
☐	aciMAC2	MAC
☐	aciIPv6Standard	IPv6 Standard
☐	aciIPv6Extended	IPv6 Extended
☐	aciARP	ARP

Delete Revert

Configuring a Standard IPv4 ACL Use the Security > ACL (Configure ACL - Add Rule - IP Standard) page to configure a Standard IPv4 ACL.

Parameters

These parameters are displayed:

- ◆ **Type** – Selects the type of ACLs to show in the Name list.
- ◆ **Name** – Shows the names of ACLs matching the selected type.
- ◆ **Action** – An ACL can contain any combination of permit or deny rules.
- ◆ **Address Type** – Specifies the source IP address. Use “Any” to include all possible addresses, “Host” to specify a specific host address in the Address field, or “IP” to specify a range of addresses with the Address and Subnet Mask fields. (Options: Any, Host, IP; Default: Any)
- ◆ **Source IP Address** – Source IP address.
- ◆ **Source Subnet Mask** – A subnet mask containing four integers from 0 to 255, each separated by a period. The mask uses 1 bits to indicate “match” and 0 bits to indicate “ignore.” The mask is bitwise ANDed with the specified source IP address, and compared with the address for each IP packet entering the port(s) to which this ACL has been assigned.
- ◆ **Time Range** – Name of a time range.

Web Interface

To add rules to an IPv4 Standard ACL:

1. Click Security, ACL.
2. Select Configure ACL from the Step list.
3. Select Add Rule from the Action list.
4. Select IP Standard from the Type list.
5. Select the name of an ACL from the Name list.
6. Specify the action (i.e., Permit or Deny).
7. Select the address type (Any, Host, or IP).
8. If you select "Host," enter a specific address. If you select "IP," enter a subnet address and the mask for an address range.
9. Click Apply.

Figure 209: Configuring a Standard IPv4 ACL

The screenshot shows the 'Security > ACL' configuration page. At the top, the breadcrumb is 'Security > ACL'. Below it, the 'Step' is '2. Configure ACL' and the 'Action' is 'Add Rule'. The 'Type' section has radio buttons for 'IP Standard' (selected), 'IP Extended', 'MAC', 'IPv6 Standard', 'IPv6 Extended', and 'ARP'. The 'Name' field is 'R&D'. The 'Action' dropdown is 'Permit'. The 'Address Type' dropdown is 'Host'. The 'Source IP Address' field is '10.1.1.21'. The 'Source Subnet Mask' field is '255.255.255.255'. The 'Time-Range' checkbox is checked, and the 'Time-Range' dropdown is '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 329](#).)
- ◆ **Source/Destination Port** – Source/destination port number for the specified protocol type. (Range: 0-65535)
- ◆ **Source/Destination Port Bit Mask** – Decimal number representing the port bits to match. (Range: 0-65535)
- ◆ **Protocol** – Specifies the protocol type to match as TCP, UDP or Others, where others indicates a specific protocol number (0-255). (Options: TCP, UDP, Others; Default: Others)

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:
 - **ToS** – Type of Service level. (Range: 0-15)
 - **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 210: Configuring an Extended IPv4 ACL

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

Configuring a Standard IPv6 ACL Use the Security > ACL (Configure ACL - Add Rule - IPv6 Standard) page to configure a Standard IPv6 ACL.

Parameters

These parameters are displayed:

- ◆ **Type** – Selects the type of ACLs to show in the Name list.
- ◆ **Name** – Shows the names of ACLs matching the selected type.
- ◆ **Action** – An ACL can contain any combination of permit or deny rules.
- ◆ **Source Address Type** – Specifies the source IP address. Use “Any” to include all possible addresses, “Host” to specify a specific host address in the Address field, or “IPv6-Prefix” to specify a range of addresses. (Options: Any, Host, IPv6-Prefix; Default: Any)
- ◆ **Source IPv6 Address** – An IPv6 source address or network class. The address must be formatted according to RFC 2373 “IPv6 Addressing Architecture,” using 8 colon-separated 16-bit hexadecimal values. One double colon may be used in the address to indicate the appropriate number of zeros required to fill the undefined fields.
- ◆ **Source Prefix-Length** – A decimal value indicating how many contiguous bits (from the left) of the address comprise the prefix (i.e., the network portion of the address). (Range: 0-128 bits)
- ◆ **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 211: Configuring a Standard IPv6 ACL

The screenshot shows the 'Security > ACL' configuration page. At the top, the breadcrumb 'Security > ACL' is displayed. Below it, the 'Step' is set to '2. Configure ACL' and the 'Action' is 'Add Rule'. The 'Type' section has radio buttons for 'IP Standard', 'IP Extended', 'MAC', 'IPv6 Standard' (which is selected), 'IPv6 Extended', and 'ARP'. The 'Name' field is set to 'R&D#6S'. The 'Action' dropdown is set to 'Permit'. The 'Source Address Type' dropdown is set to 'Host'. The 'Source IPv6 Address' field contains '2009:DB9:2229::79'. The 'Source Prefix Length (0-128)' field contains '128'. There is a checkbox for 'Time-Range' which is unchecked, and a dropdown for 'R&D' is visible. At the bottom right, there are 'Apply' and 'Revert' buttons.

Configuring an Extended IPv6 ACL Use the Security > ACL (Configure ACL - Add Rule - IPv6 Extended) page to configure an Extended IPv6 ACL.

Parameters

These parameters are displayed:

- ◆ **Type** – Selects the type of ACLs to show in the Name list.
- ◆ **Name** – Shows the names of ACLs matching the selected type.

- ◆ **Action** – An ACL can contain any combination of permit or deny rules.
- ◆ **Destination Address Type** – Specifies the destination IP address type. Use “Any” to include all possible addresses, “Host” to specify a specific host address in the Address field, or “IPv6-Prefix” to specify a range of addresses. (Options: Any, Host, IPv6-Prefix; Default: Any)
- ◆ **Destination IPv6 Address** – An IPv6 address or network class. The address must be formatted according to RFC 2373 “IPv6 Addressing Architecture,” using 8 colon-separated 16-bit hexadecimal values. One double colon may be used in the address to indicate the appropriate number of zeros required to fill the undefined fields.
- ◆ **Destination Prefix Length** – A decimal value indicating how many contiguous bits (from the left) of the address comprise the prefix; i.e., the network portion of the address. (Range: 0-128 bits)

Web Interface

To add rules to an Extended IPv6 ACL:

1. Click Security, ACL.
2. Select Configure ACL from the Step list.
3. Select Add Rule from the Action list.
4. Select IPv6 Extended from the Type list.
5. Select the name of an ACL from the Name list.
6. Specify the action (i.e., Permit or Deny).
7. Select the address type (Any, Host or IPv6-prefix).
8. If you select “Host,” enter a specific address. If you select “IPv6-prefix,” enter a subnet address and prefix length.
9. Click Apply.

Figure 212: Configuring an Extended IPv6 ACL

The screenshot shows the 'Security > ACL' configuration page. At the top, there's a breadcrumb 'Security > ACL'. Below it, a 'Step:' dropdown is set to '1. Configure ACL' and an 'Action:' dropdown is set to 'Add Rule'. The 'Type' section has five radio buttons: 'IP Standard', 'IP Extended', 'MAC', 'IPv6 Standard', and 'IPv6 Extended'. The 'IPv6 Extended' option is selected. The 'Name' field contains 'ipv6e'. The 'Action' dropdown is set to 'Permit'. The 'Destination Address Type' dropdown is set to 'IPv6-Prefix'. The 'Destination IPv6 Address' field contains '2009:db9:2229::78'. The 'Destination Prefix Length (0-128)' field contains '8'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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, 4095 meaning exact match)

- ◆ **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.)
- ◆ **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 213: Configuring a MAC ACL

The screenshot shows the 'Security > ACL' configuration page. The 'Step' is '2. Configure ACL' and the 'Action' is 'Add Rule'. The 'Type' is set to 'MAC'. The 'Name' is 'mac'. The 'Action' is 'Permit'. The 'Source Address Type' is 'Any', 'Source MAC Address' is '00-00-00-00-00-00', 'Source Bit Mask' is '00-00-00-00-00-00', 'Destination Address Type' is 'Any', 'Destination MAC Address' is '00-00-00-00-00-00', and 'Destination Bit Mask' is '00-00-00-00-00-00'. The 'Packet Format' is 'Any'. The 'VID (1-4094)' is empty, 'VID Bit Mask (0-4095)' is empty, 'CoS (0-7)' is empty, and 'CoS Bit Mask (0-7)' is empty. The 'Time-Range' is 'R&D'. The 'Ethernet Type' is '(0000-FFFF, hexadecimal value)' and 'Ethernet Type Bit Mask' is '(0000-FFFF, hexadecimal value)'. The 'Apply' and 'Revert' buttons are at the bottom.

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 386](#)).

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 329](#).)

- ◆ **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 214: Configuring a ARP ACL

The screenshot shows the 'Security > ACL' configuration page. At the top, the 'Step' is '2. Configure ACL' and the 'Action' is 'Add Rule'. The 'Type' is set to 'ARP'. The 'Name' is 'R&F#7ARP'. The 'Action' is 'Permit'. The 'Packet Type' is 'IP'. The 'Source IP Address Type' is 'Any', 'Source IP Address' is '0.0.0.0', and 'Source IP Subnet Mask' is '0.0.0.0'. The 'Destination IP Address Type' is 'Any', 'Destination IP Address' is '0.0.0.0', and 'Destination IP Subnet Mask' is '0.0.0.0'. The 'Source MAC Address Type' is 'Any', 'Source MAC Address' is '00-00-00-00-00-00', and 'Source MAC Bit Mask' is '00-00-00-00-00-00'. The 'Destination MAC Address Type' is 'Any', 'Destination MAC Address' is '00-00-00-00-00-00', and 'Destination MAC Bit Mask' is '00-00-00-00-00-00'. 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 from the Action list.
3. Select IP or MAC from the Type list.
4. Select a port.

5. Select the name of an ACL from the ACL list.
6. Select a time range during which the ACL is active.
7. Enable the counter for ACL statistics if required.
8. Click Apply.

Figure 215: Binding a Port to an ACL

The screenshot shows the 'Security > ACL' configuration page, specifically 'Step: 3. Configure Interface'. The 'Type' section has radio buttons for IP (selected), MAC, and IPv6. The 'Port' is set to 1. Under the 'Ingress' section, the 'ACL' checkbox is checked with a dropdown set to 'ips', 'Time-Range' is unchecked with a dropdown set to 'R&D', and 'Counter' is unchecked. Under the 'Egress' section, 'ACL' is unchecked with a dropdown set to 'ips', 'Time-Range' is unchecked with a dropdown set to 'R&D', and 'Counter' is unchecked. 'Apply' and 'Revert' buttons are at the bottom right.

Showing ACL Hardware Counters Use the Security > ACL > (Configure Interface - Show Hardware Counters) page to show statistics for ACL hardware counters.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28/52)
- ◆ **Type** – Selects the type of ACL.
- ◆ **Direction** – Displays statistics for ingress or egress traffic.
- ◆ **Query** – Displays statistics for selected criteria.
- ◆ **ACL Name** – The ACL bound this port.
- ◆ **Action** – Shows if action is to permit or deny specified packets.
- ◆ *Rules* – Shows the rules for the ACL bound to this port.
- ◆ **Time-Range** – Name of a time range.

- ◆ **Hit** – Shows the number of packets matching this ACL.
- ◆ **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 216: Showing ACL Statistics

The screenshot shows the 'Security > ACL' web interface. At the top, there's a breadcrumb 'Security > ACL' and a help icon. Below that, there are two dropdown menus: 'Steps: 3. Configure Interface' and 'Action: Show Hardware Counter'. The main configuration area has three dropdown menus: 'Port: 1', 'Type: IP Standard', and 'Direction: Ingress'. A 'Query' button is located to the right of these settings. Below the configuration area, it says 'ACL Name: rd' and 'Total: 1'. At the bottom, there is a table with the following data:

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 217: Creating an IP Address Filter for Management Access

Security > IP Filter

Action: Add

Mode ☒ Web ☐ SNMP ☐ Telnet ☐ All

Start IP Address

End IP Address

Apply Revert

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

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

Figure 218: Showing IP Addresses Authorized for Management Access

Security > IP Filter

Action: Show

Mode ☐ Web ☒ SNMP ☐ Telnet ☐ All

SNMP IP Filter List Total: 1

	Start IP Address	End IP Address
☐	10.1.2.3	10.1.2.3

Delete Revert

Configuring Port Security

Use the Security > Port Security page to configure the maximum number of device MAC addresses that can be learned by a switch port, stored in the address table, and authorized to access the network.

When port security is enabled on a port, the switch stops learning new MAC addresses on the specified port when it has reached a configured maximum number. Only incoming traffic with source addresses already stored in the address table will be authorized to access the network through that port. If a device with an unauthorized MAC address attempts to use the switch port, the intrusion will be detected and the switch can automatically take action by disabling the port and sending a trap message.

Command Usage

- ◆ The default maximum number of MAC addresses allowed on a secure port is zero (that is, disabled). To use port security, you must configure the maximum number of addresses allowed on a port.

- ◆ To configure the maximum number of address entries which can be learned on a port, 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 201](#)).

- ◆ 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 106](#)).
- ◆ 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, 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 310](#).
- ◆ **MAC Filter ID** – The identifier for a MAC address filter.
- ◆ **Last Intrusion MAC** – The last unauthorized MAC address detected.
- ◆ **Last Time Detected Intrusion MAC** – The last time an unauthorized MAC address was detected.

Web Interface

To configure port security:

1. Click Security, Port Security.
2. Mark the check box in the Security Status column to enable security, set the action to take when an invalid address is detected on a port, and set the maximum number of MAC addresses allowed on the port.
3. Click Apply.

Figure 219: Configuring Port Security

Security > Port Security									
Port Security List Total: 26									
Port	Security Status	Port Status	Action	Max MAC Count (0-1024)	Current MAC Count	MAC Filter	MAC Filter ID	Last Intrusion MAC	Last Time Detected Intrusion MAC
1	☒ Enabled	Secure/Down	Trap and Shutdown	0	1	Disabled	0	NA	NA
2	☒ Enabled	Secure/Down	Trap	0	0	Disabled	1	NA	NA
3	☐ Enabled	Secure/Down	None	0	0	Disabled	0	NA	NA


Configuring 802.1X Port Authentication

Network switches can provide open and easy access to network resources by simply attaching a client PC. Although this automatic configuration and access is a desirable feature, it also allows unauthorized personnel to easily intrude and possibly gain access to sensitive network data.

The IEEE 802.1X (dot1X) standard defines a port-based access control procedure that prevents unauthorized access to a network by requiring users to first submit credentials for authentication. Access to all switch ports in a network can be centrally controlled from a server, which means that authorized users can use the same credentials for authentication from any point within the network.

This switch uses the Extensible Authentication Protocol over LANs (EAPOL) to exchange authentication protocol messages with the client, and a remote RADIUS authentication server to verify user identity and access rights. When a client (i.e., Supplicant) connects to a switch port, the switch (i.e., Authenticator) responds with an EAPOL identity request. The client provides its identity (such as a user name) in an EAPOL response to the switch, which it forwards to the RADIUS server. The RADIUS server verifies the client identity and sends an access challenge back to the client. The EAP packet from the RADIUS server contains not only the challenge, but the authentication method to be used. The client can reject the authentication method and request another, depending on the configuration of the client software and the RADIUS server. The encryption method used to pass authentication messages can be MD5 (Message-Digest 5), TLS (Transport Layer Security), PEAP (Protected Extensible Authentication Protocol), or TTLS (Tunneled Transport Layer Security). The client responds to the appropriate method with its credentials, such as a password or certificate. The RADIUS server verifies the client credentials and responds with an accept or reject packet. If authentication is successful, the switch allows the client to access the network. Otherwise, non-EAP traffic on the port is blocked or assigned to a guest VLAN based on the “intrusion-action” setting. In “multi-host” mode, only one host connected to a port needs to pass authentication for all other hosts to be granted network access. Similarly, a port can become unauthorized for all hosts if one attached host fails re-authentication or sends an EAPOL logoff message.

Figure 220: Configuring Port Authentication



The operation of 802.1X on the switch requires the following:

- ◆ The switch must have an IP address assigned.
- ◆ RADIUS authentication must be enabled on the switch and the IP address of the RADIUS server specified.
- ◆ 802.1X must be enabled globally for the switch.
- ◆ Each switch port that will be used must be set to dot1X "Auto" mode.
- ◆ Each client that needs to be authenticated must have dot1X client software installed and properly configured.
- ◆ The RADIUS server and 802.1X client support EAP. (The switch only supports EAPOL in order to pass the EAP packets from the server to the client.)
- ◆ The RADIUS server and client also have to support the same EAP authentication type – MD5, PEAP, TLS, or TTLS. (Native support for these encryption methods is provided in Windows 7, 8 and 10.)

Configuring 802.1X Global Settings

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

Parameters

These parameters are displayed:

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

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

When this device is functioning as an edge switch but does not require any attached clients to be authenticated, **EAPOL Pass Through** can be disabled to discard unnecessary EAPOL traffic.
- ◆ **Default** – Sets all configurable 802.1X global and port settings to their default values.

Web Interface

To configure global settings for 802.1X:

1. Click Security, Port Authentication.
2. Select Configure Global from the Step list.
3. Enable 802.1X globally for the switch, and configure EAPOL Pass Through if required.
4. Click Apply

Figure 221: Configuring Global Settings for 802.1X Port Authentication

Security > Port Authentication

Step: 1. Configure Global

System Authentication Control ☐ Enabled

EAPOL Pass Through ☐ Enabled

Apply Revert

Default [Click this button to set 802.1X global/port settings to default values.](#)

Configuring Port Authenticator Settings for 802.1X

Use the Security > Port Authentication (Configure Interface) 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 286.](#))
- ◆ **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 163](#)) and mapped on each port (See [“Configuring Network Access for Ports” on page 307](#)).

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 222: Configuring Interface Settings for 802.1X Port Authenticator

Security > Port Authentication

Step: 2. Configure Interface

Port: 1

Status: Disabled

Authorized: Yes

Control Mode: Auto

Operation Mode: Single-Host

Max MAC Count (1-1024): 5

Max Request (1-10): 2

Quiet Period (1-65535): 60 sec

Tx Period (1-65535): 30 sec

Supplicant Timeout (1-65535): 30 sec

Server Timeout: 10 sec

Re-authentication Status: ☐ Enabled

Re-authentication Period (1-65535): 3600 sec

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

Intrusion Action: Block Traffic

Intrusion Action: Block Traffic

Supplicant List Total: 1

Supplicant	Authenticator PAE State Machine			Backend State Machine			Reauthentication State Machine
	State	Reauth Count	Current Identifier	State	Request Count	Identifier (Server)	State
00-00-00-00-00-00	Initialize	0	0	Initialize	0	0	Initialize

Apply Revert

Displaying 802.1X Statistics Use the Security > Port Authentication (Show Statistics) page to display statistics for dot1x protocol exchanges for any port.

Parameters

These parameters are displayed:

Table 23: 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.

Table 23: 802.1X Statistics (Continued)

Parameter	Description
Rx Last EAPOLSrc	The source MAC address carried in the most recent EAPOL frame received by this Authenticator.
Rx EAP Resp/Id	The number of EAP Resp/Id frames that have been received by this Authenticator.
Rx EAP Resp/Oth	The number of valid EAP Response frames (other than Resp/Id frames) that have been received by this Authenticator.
Rx EAP LenError	The number of EAPOL frames that have been received by this Authenticator in which the Packet Body Length field is invalid.
Tx EAP Req/Id	The number of EAP Req/Id frames that have been transmitted by this Authenticator.
Tx EAP Req/Oth	The number of EAP Request frames (other than Rq/Id frames) that have been transmitted by this Authenticator.
Tx EAPOL Total	The number of EAPOL frames of any type that have been transmitted by this Authenticator.
<i>Supplicant</i>	
Rx EAPOL Invalid	The number of EAPOL frames that have been received by this Supplicant in which the frame type is not recognized.
Rx EAPOL Total	The number of valid EAPOL frames of any type that have been received by this Supplicant.
Rx Last EAPOLVer	The protocol version number carried in the most recent EAPOL frame received by this Supplicant.
Rx Last EAPOLSrc	The source MAC address carried in the most recent EAPOL frame received by this Supplicant.
Rx EAP Resp/Id	The number of EAP Resp/Id frames that have been received by this Supplicant.
Rx EAP Resp/Oth	The number of valid EAP Response frames (other than Resp/Id frames) that have been received by this Supplicant.
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.

Figure 223: Showing Statistics for 802.1X Port Authenticator

The screenshot shows the 'Security > Port Authentication' web interface. At the top, there is a breadcrumb trail 'Security > Port Authentication'. Below it, a 'Step:' dropdown menu is set to '3. Show Statistics'. Underneath, a 'Port' dropdown menu is set to '1'. The main section is titled 'Port Authentication Authenticator Statistics'. It contains a table with two columns of statistics. The first column lists statistics for 'Rx' (receive) and the second column lists statistics for 'Tx' (transmit). The table data is as follows:

Statistic	Value	Statistic	Value
Rx EAPOL Start	1	Rx EAP Resp/Id	0
Rx EAPOL Logoff	0	Rx EAP Resp/Oth	0
Rx EAPOL Invalid	0	Rx EAP LenError	0
Rx EAPOL Total	1	Tx EAP Req/Id	0
Rx Last EAPOLVer	2	Tx EAP Req/Oth	0
Rx Last EAPOLSrc	70-72-CF-32-DD-FD	Tx EAPOL Total	0

At the bottom right of the table, there is a 'Refresh' button.

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:

- ◆ **LAND Attack** – Configures the switch to protect against DoS LAND (Local Area Network Denial) attacks in which hackers send spoofed-IP packets where the source and destination address are the same, thereby causing the target to reply to itself continuously. (Default: Enabled)
- ◆ **TCP Null Scan** – A TCP NULL scan message is used to identify listening TCP ports. The scan uses a series of strangely configured TCP packets which contain a sequence number of 0 and no flags. If the target's TCP port is closed, the target replies with a TCP RST (reset) packet. If the target TCP port is open, it simply discards the TCP NULL scan. (Default: Enabled)


- ◆ **TCP-SYN/FIN Scan** – A TCP SYN/FIN scan message is used to identify listening TCP ports. The scan uses a series of strangely configured TCP packets which contain SYN (synchronize) and FIN (finish) flags. If the target's TCP port is closed, the target replies with a TCP RST (reset) packet. If the target TCP port is open, it simply discards the TCP SYN FIN scan. (Default: Enabled)
- ◆ **TCP Xmas Scan** – A so-called TCP XMAS scan message is used to identify listening TCP ports. This scan uses a series of strangely configured TCP packets which contain a sequence number of 0 and the URG, PSH and FIN flags. If the target's TCP port is closed, the target replies with a TCP RST packet. If the target TCP port is open, it simply discards the TCP XMAS scan. (Default: Enabled)

Web Interface

To protect against DoS attacks:

1. Click Security, DoS Protection.
2. Enable protection for specific DoS attacks.
3. Click Apply

Figure 224: Protecting Against DoS Attacks



DHCPv4 Snooping

The addresses assigned to DHCPv4 clients on insecure ports can be carefully controlled using the dynamic bindings registered with DHCPv4 Snooping (or using the static bindings configured with IPv4 Source Guard). DHCPv4 snooping allows a switch to protect a network from rogue DHCPv4 servers or other devices which send port-related information to a DHCPv4 server. This information can be useful in tracking an IPv4 address back to a physical port.

Command Usage

DHCP Snooping Process

- ◆ Network traffic may be disrupted when malicious DHCP messages are received from an outside source. DHCP snooping is used to filter DHCP messages received on a non-secure interface from outside the network or fire wall. When DHCP snooping is enabled globally and enabled on a VLAN interface, DHCP

messages received on an untrusted interface from a device not listed in the DHCP snooping table will be dropped.

- ◆ Table entries are only learned for trusted interfaces. An entry is added or removed dynamically to the DHCP snooping table when a client receives or releases an IP address from a DHCP server. Each entry includes a MAC address, IP address, lease time, VLAN identifier, and port identifier.
- ◆ The rate limit for the number of DHCP messages that can be processed by the switch is 100 packets per second. Any DHCP packets in excess of this limit are dropped.
- ◆ When DHCP snooping is enabled, DHCP messages entering an untrusted interface are filtered based upon dynamic entries learned via DHCP snooping.
- ◆ Filtering rules are implemented as follows:
 - If the global DHCP snooping is disabled, all DHCP packets are forwarded.
 - If DHCP snooping is enabled globally, and also enabled on the VLAN where the DHCP packet is received, all DHCP packets are forwarded for a *trusted* port. If the received packet is a DHCP ACK message, a dynamic DHCP snooping entry is also added to the binding table.
 - If DHCP snooping is enabled globally, and also enabled on the VLAN where the DHCP packet is received, but the port is *not trusted*, it is processed as follows:
 - If the DHCP packet is a reply packet from a DHCP server (including OFFER, ACK or NAK messages), the packet is dropped.
 - If the DHCP packet is from a client, such as a DECLINE or RELEASE message, the switch forwards the packet only if the corresponding entry is found in the binding table.
 - If the DHCP packet is from a client, such as a DISCOVER, REQUEST, INFORM, DECLINE or RELEASE message, the packet is forwarded if MAC address verification is disabled. However, if MAC address verification is enabled, then the packet will only be forwarded if the client's hardware address stored in the DHCP packet is the same as the source MAC address in the Ethernet header.
 - If the DHCP packet is not a recognizable type, it is dropped.
 - If a DHCP packet from a client passes the filtering criteria above, it will only be forwarded to trusted ports in the same VLAN.
 - If a DHCP packet is from server is received on a trusted port, it will be forwarded to both trusted and untrusted ports in the same VLAN.
 - If the DHCP snooping is globally disabled, all dynamic bindings are removed from the binding table.

- *Additional considerations when the switch itself is a DHCP client* – The port(s) through which the switch submits a client request to the DHCP server must be configured as trusted. Note that the switch will not add a dynamic entry for itself to the binding table when it receives an ACK message from a DHCP server. Also, when the switch sends out DHCP client packets for itself, no filtering takes place. However, when the switch receives any messages from a DHCP server, any packets received from untrusted ports are dropped.

DHCP Snooping Option 82

- ◆ DHCP provides a relay mechanism for sending information about its DHCP clients or the relay agent itself to the DHCP server. Also known as DHCP Option 82, it allows compatible DHCP servers to use the information when assigning IP addresses, or to set other services or policies for clients. It is also an effective tool in preventing malicious network attacks from attached clients on DHCP services, such as IP Spoofing, Client Identifier Spoofing, MAC Address Spoofing, and Address Exhaustion.
- ◆ DHCP Snooping must be enabled for Option 82 information to be inserted into request packets.
- ◆ When the DHCP Snooping Information Option 82 is enabled, the requesting client (or an intermediate relay agent that has used the information fields to describe itself) can be identified in the DHCP request packets forwarded by the switch and in reply packets sent back from the DHCP server. This information may specify the MAC address or IP address of the requesting device (that is, the switch in this context).

By default, the switch also fills in the Option 82 circuit-id field with information indicating the local interface over which the switch received the DHCP client request, including the port and VLAN ID. This allows DHCP client-server exchange messages to be forwarded between the server and client without having to flood them to the entire VLAN.

- ◆ If DHCP Snooping Information Option 82 is enabled on the switch, information may be inserted into a DHCP request packet received over any VLAN (depending on DHCP snooping filtering rules). The information inserted into the relayed packets includes the circuit-id and remote-id, as well as the gateway Internet address.
- ◆ When the switch receives DHCP packets from clients that already include DHCP Option 82 information, the switch can be configured to set the action policy for these packets. The switch can either drop the DHCP packets, keep the existing information, or replace it with the switch's relay information.

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

Parameters

These parameters are displayed:

General

- ◆ **DHCP Snooping Status** – Enables DHCP snooping globally. (Default: Disabled)
- ◆ **DHCP Snooping MAC-Address Verification** – Enables or disables MAC address verification. If the source MAC address in the Ethernet header of the packet is not same as the client's hardware address in the DHCP packet, the packet is dropped. (Default: Enabled)
- ◆ **DHCP Snooping Rate Limit** – Sets the maximum number of DHCP packets that can be trapped by the switch for DHCP snooping. (Range: 1-2048 packets/second)

Information

- ◆ **DHCP Snooping Information Option Status** – Enables or disables DHCP Option 82 information relay. (Default: Disabled)
- ◆ **DHCP Snooping Information Option Sub-option Format** – Enables or disables use of sub-type and sub-length fields in circuit-ID (CID) and remote-ID (RID) in Option 82 information. (Default: Enabled)
- ◆ **DHCP Snooping Information Option Remote ID** – Specifies the MAC address, IP address, or arbitrary identifier of the requesting device (i.e., the switch in this context).
 - **MAC Address** – Inserts a MAC address in the remote ID sub-option for the DHCP snooping agent (i.e., the MAC address of the switch's CPU). This attribute can be encoded in Hexadecimal or ASCII.
 - **IP Address** – Inserts an IP address in the remote ID sub-option for the DHCP snooping agent (i.e., the IP address of the management interface). This attribute can be encoded in Hexadecimal or ASCII.
 - *string* – An arbitrary string inserted into the remote identifier field. (Range: 1-32 characters)
- ◆ **DHCP Snooping Information Option Policy** – Specifies how to handle DHCP client request packets which already contain Option 82 information.
 - **Drop** – Drops the client's request packet instead of relaying it.
 - **Keep** – Retains the Option 82 information in the client request, and forwards the packets to trusted ports.
 - **Replace** – Replaces the Option 82 information circuit-id and remote-id fields in the client's request with information about the relay agent itself,

inserts the relay agent's address (when DHCP snooping is enabled), and forwards the packets to trusted ports. (This is the default policy.)

Web Interface

To configure global settings for DHCP Snooping:

1. Click IP Service, DHCP, Snooping.
2. Select Configure Global from the Step list.
3. Select the required options for the general DHCP snooping process and for the DHCP snooping information options.
4. Click Apply

Figure 225: Configuring Global Settings for DHCP Snooping

The screenshot shows the 'Security > DHCP Snooping' configuration page. At the top, there's a breadcrumb trail 'Security > DHCP Snooping' and a 'Step:' dropdown menu set to '1. Configure Global'. Below this, the 'General' section contains three settings: 'DHCP Snooping Status' (checkbox, disabled), 'DHCP Snooping MAC-Address Verification' (checkbox, checked/enabled), and 'DHCP Snooping Rate Limit' (input field with '0' and '(pkt/sec)' label). The 'Information' section contains four settings: 'DHCP Snooping Information Option Status' (checkbox, disabled), 'DHCP Snooping Information Option Sub-option Format' (dropdown menu, 'Extra Subtype Included'), 'DHCP Snooping Information Option Remote ID' (dropdown menu, 'MAC Address (Hex Encoded)' with an adjacent input field), and 'DHCP Snooping Information Option Policy' (dropdown menu, 'Replace'). At the bottom right of the form are 'Apply' and 'Revert' buttons.

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

Command Usage

- ◆ When DHCP snooping is enabled globally on the switch, and enabled on the specified VLAN, DHCP packet filtering will be performed on any untrusted ports within the VLAN.
- ◆ When the DHCP snooping is globally disabled, DHCP snooping can still be configured for specific VLANs, but the changes will not take effect until DHCP snooping is globally re-enabled.
- ◆ When DHCP snooping is globally enabled, and DHCP snooping is then disabled on a VLAN, all dynamic bindings learned for this VLAN are removed from the binding table.

Parameters

These parameters are displayed:

- ◆ **VLAN** – ID of a configured VLAN. (Range: 1-4094)
- ◆ **DHCP Snooping Status** – Enables or disables DHCP snooping for the selected VLAN. When DHCP snooping is enabled globally on the switch, and enabled on the specified VLAN, DHCP packet filtering will be performed on any untrusted ports within the VLAN. (Default: Disabled)

Web Interface

To configure global settings for DHCP Snooping:

1. Click IP Service, DHCP, Snooping.
2. Select Configure VLAN from the Step list.
3. Enable DHCP Snooping on any existing VLAN.
4. Click Apply

Figure 226: Configuring DHCP Snooping on a VLAN

The screenshot shows a web interface for configuring DHCP Snooping. The breadcrumb trail is 'Security > DHCP Snooping'. Below this, there is a 'Step:' dropdown menu currently showing '2. Configure VLAN'. Underneath, there are two configuration fields: 'VLAN' with a dropdown menu showing '1', and 'DHCP Snooping Status' with a checked checkbox labeled 'Enabled'. At the bottom right of the configuration area, there are two buttons: 'Apply' and 'Revert'.

Configuring Interfaces for DHCP Snooping Use the IP Service > DHCP > Snooping (Configure Interface) page to configure switch interfaces as trusted or untrusted.

Command Usage

- ◆ A trusted interface is an interface that is configured to receive only messages from within the network. An untrusted interface is an interface that is configured to receive messages from outside the network or fire wall.
- ◆ When DHCP snooping is enabled both globally and on a VLAN, DHCP packet filtering will be performed on any untrusted interface within the VLAN.
- ◆ When an untrusted interface is changed to a trusted port, all the dynamic DHCP snooping bindings associated with this interface is removed.
- ◆ Set all interfaces connected to DHCP servers within the local network or fire wall to trusted state. Set all other interfaces outside the local network or fire wall to untrusted state.

Parameters

These parameters are displayed:

- ◆ **Interface** – Port or trunk identifier.
- ◆ **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: 1-5; Default: 5)
- ◆ **Circuit ID** – Specifies DHCP Option 82 circuit ID suboption information.
 - **Mode** – Specifies the default string “VLAN-Unit-Port” or an arbitrary string. (Default: VLAN-Unit-Port)
 - **Value** – An arbitrary string inserted into the circuit identifier field. (Range: 1-32 characters)

Web Interface

To configure global settings for DHCP Snooping:

1. Click IP Service, DHCP, Snooping.
2. Select Configure Interface from the Step list.
3. Set any ports within the local network or firewall to trusted.
4. Specify the mode used for sending circuit ID information, and an arbitrary string if required.
5. Click Apply

Figure 227: Configuring the Port Mode for DHCP Snooping

Security > DHCP Snooping

Step: 3. Configure Interface

Interface: ☒ Port ☐ Trunk

DHCP Snooping Port List Total: 28

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

Displaying DHCP Snooping Binding Information

Use the IP Service > DHCP > Snooping (Show Information) page to display entries in the binding table.

Parameters

These parameters are displayed:

- ◆ **Interface**
 - Port identifier. (Range: 1-28/52)
 - Trunk identifier. (Range: 1-26)
- ◆ **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-DHCPsnp** – Statically configured.
- ◆ **VLAN** – VLAN to which this entry is bound.
- ◆ **Interface** – Port or trunk to which this entry is bound.
- ◆ **Clear** – Removes all dynamically learned snooping entries from RAM.
- ◆ **Store to Flash** – 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 from Flash** – 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 228: Displaying the Binding Table for DHCP Snooping

IP Service > DHCP > Snooping

Step: 4. Show Information

DHCP Snooping Binding List Total: 6

MAC Address	IP Address	Lease Time (seconds)	Type	VLAN	Interface
00-10-B5-F4-00-01	10.2.44.96	5	DHCP-Snooping	1	Trunk 1
00-10-B5-F4-00-02	10.3.44.96	15	Static-DHCPSPNP	1	Unit 1 / Port 2
00-10-B5-F4-00-03	10.4.44.96	25	DHCP-Snooping	1	Unit 1 / Port 3
00-10-B5-F4-00-04	10.5.44.96	10	Static-DHCPSPNP	1	Trunk 4
00-10-B5-F4-00-05	10.6.44.96	10	DHCP-Snooping	1	Unit 1 / Port 5
00-10-B5-F4-00-06	10.7.44.96	5	Static-DHCPSPNP	1	Unit 1 / Port 6

Store Click the button to Store DHCP Snooping binding entries to flash.

Clear Click the button to Clear DHCP Snooping binding entries from flash.

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 an unsecure 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 IP Service > DHCPv6 > Snooping (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 IP Service, DHCPv6, Snooping.
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 229: 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 IP Service > DHCPv6 > Snooping (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 IP Service, DHCPv6, Snooping.
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 230: 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 IP Service, DHCPv6, Snooping.
2. Select Configure VLAN from the Step list.
3. Select Show from the Action list.

Figure 231: 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 IP Service > 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 IP Service, DHCPv6, Snooping.
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 232: Configuring the Trust Sate for DHCPv6 Snooping



Interface	Trust Status	Max Binding (1 - 5)	Current Binding
Eth 1/1	☒ Enabled	3	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 IP Service > DHCPv6 > Snooping (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 IP Service, DHCPv6, Snooping.
2. Select Show Information from the Step list.
3. Click Binding.
4. Use the Clear function if required.

Figure 233: Displaying the Binding Table for DHCPv6 Snooping



	Link-layer Address	IPv6 Address	Lifetime	VLAN	Interface	Type
☐	00-13-49-AA-39-26	2001:b021:1435:5612:ab3c:6792:a452:6712	2591998	1	Eth 1/5	NA
☐	00-12-CF-01-02-03	2001:b000::1	2591912	1	Eth 1/3	NA
☐	00-12-CF-01-02-04	2001:b000::3	2591555	1	Trunk 2	TA

Clear

Displaying DHCPv6 Snooping Statistics Use the IP Service > DHCPv6 > Snooping (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 IP Service, DHCPv6, Snooping.
2. Select Show Information from the Step list.
3. Click Statistics.
4. Use the Clear function if required.

Figure 234: Displaying Statistics for DHCPv6 Snooping

Security > DHCP Snooping6

Step: 4. Show Information ▼

☐ Binding ☒ Statistics

DHCPv6 Snooping Statistics

Client Packet: Solicit, Request, Confirm, Renew, Rebind, Decline, Release, Information-request

Server Packet: Advertise, Reply, Reconfigure

Relay Packet: Relay-forward, Relay-reply

State	Client	Server	Relay	Total
Received	0	0	0	0
Sent	0	0	0	0
Dropped	0	0	0	0

Clear

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 [“DHCPv4 Snooping” on page 356](#)). IPv4 source guard can be used to prevent traffic attacks caused when a host tries to use the IPv4 address of a neighbor to access the network. This section describes how to configure IPv4 Source Guard.

Configuring Ports for IPv4 Source Guard

Use the Security > IP Source Guard > Port Configuration page to set the filtering type based on source IPv4 address, or source IPv4 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.

IPv4 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 IPv4 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 [“DHCPv4 Snooping” on page 356](#)), 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 359](#)), IPv4 source guard will check the VLAN ID, source IP address, port number, and source MAC address (for the SIP-MAC option). If a matching entry is found in the binding table and

the entry type is static IPv4 source guard binding, the packet will be forwarded.

- If DHCP snooping is enabled, IPv4 source guard will check the VLAN ID, source IP address, port number, and source MAC address (for the SIP-MAC option). If a matching entry is found in the binding table and the entry type is static IPv4 source guard binding, or dynamic DHCP snooping binding, the packet will be forwarded.
- If IPv4 source guard is enabled on an interface for which IP source bindings have not yet been configured (neither by static configuration in the IPv4 source guard binding table nor dynamically learned from DHCP snooping), the switch will drop all IP traffic on that port, except for DHCP packets allowed by DHCP snooping.
- ◆ An entry with same MAC address and a different VLAN ID cannot be added to the binding table.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28/52)
- ◆ **Filter Type** – Configures the switch to filter inbound traffic based source IP address, or source IP address and corresponding MAC address. (Default: None)
 - **Disabled** – Disables IPv4 source guard filtering on the port.
 - **SIP** – Enables traffic filtering based on IPv4 addresses stored in the binding table.
 - **SIP-MAC** – Enables traffic filtering based on IPv4 addresses and corresponding MAC addresses stored in the binding table.
- ◆ **Filter Table** – Sets the source guard learning model to search for addresses in the ACL binding table or the MAC address binding table. (Default: ACL binding table)
- ◆ **Max Binding Entry** – The maximum number of entries that can be bound to an interface. (ACL Table: 1-16, default: 5; MAC Table: 1-1024, default: 1024)

This parameter sets the maximum number of address entries that can be mapped to an interface in the binding table, including both dynamic entries discovered by DHCP snooping (see [“DHCPv4 Snooping” on page 356](#)) and static entries set by IP source guard (see [“Configuring Static Bindings for IPv4 Source Guard” on page 375](#)).

The maximum binding for ACL mode restricts the number of “active” entries per port. If binding entries exceed the maximum number in IPv4 source guard, only the maximum number of binding entries will be set. Dynamic binding entries exceeding the maximum number will be created but will not be active.

The maximum binding for MAC mode restricts the number of MAC addresses learned per port. Authenticated IP traffic with different source MAC addresses cannot be learned if it would exceed this maximum number.

Web Interface

To set the IP Source Guard filter for ports:

1. Click Security, IP Source Guard, Port Configuration.
2. Set the required filtering type for each port.
3. Click Apply

Figure 235: Setting the Filter Type for IP 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	1024
2	DISABLED	ACL	5	1024
3	DISABLED	ACL	5	1024
4	DISABLED	ACL	5	1024
5	SIP	ACL	5	1024

Configuring Static Bindings for IPv4 Source Guard

Use the Security > IP Source Guard > Static Binding (Configure ACL Table and Configure MAC Table) pages to bind a static address to a port. Table entries include a MAC address, IP address, lease time, entry type (Static, Dynamic), VLAN identifier, and port identifier. All static entries are configured with an infinite lease time, which is indicated with a value of zero in the table.

Command Usage

- ◆ Table entries include a MAC address, IP address, lease time, entry type (Static-IP-SG-Binding, Dynamic-DHCP-Binding), VLAN identifier, and port identifier.
- ◆ Static addresses entered in the source guard binding table are automatically configured with an infinite lease time.
- ◆ When source guard is enabled, traffic is filtered based upon dynamic entries learned via DHCP snooping, or static addresses configured in the source guard binding table.
- ◆ An entry with same MAC address and a different VLAN ID cannot be added to the binding table.
- ◆ Static bindings are processed as follows:
 - A valid static IP source guard entry will be added to the binding table in ACL mode if one of the following conditions is true:

- If there is no entry with the same VLAN ID and MAC address, a new entry is added to the binding table using the type “static IP source guard binding.”
- If there is an entry with the same VLAN ID and MAC address, and the type of entry is static IP source guard binding, then the new entry will replace the old one.
- If there is an entry with the same VLAN ID and MAC address, and the type of the entry is dynamic DHCP snooping binding, then the new entry will replace the old one and the entry type will be changed to static IP source guard binding.
- A valid static IP source guard entry will be added to the binding table in MAC mode if one of the following conditions are true:
 - If there is no binding entry with the same IP address and MAC address, a new entry will be added to the binding table using the type of static IP source guard binding entry.
 - If there is a binding entry with same IP address and MAC address, then the new entry shall replace the old one.
- Only unicast addresses are accepted for static bindings.

Parameters

These parameters are displayed:

Add – Configure ACL Table

- ◆ **Port** – The port to which a static entry is bound.
- ◆ **VLAN** – ID of a configured VLAN (Range: 1-4094)
- ◆ **MAC Address** – A valid unicast MAC address.
- ◆ **IP Address** – A valid unicast IP address, including classful types A, B or C.

Add – Configure MAC Table

- ◆ **MAC Address** – A valid unicast MAC address.
- ◆ **VLAN** – ID of a configured VLAN or a range of VLANs. (Range: 1-4094)
- ◆ **IP Address** – A valid unicast IP address, including classful types A, B or C.
- ◆ **Port** – The port to which a static entry is bound. Specify a physical port number or list of port numbers. Separate nonconsecutive port numbers with a comma and no spaces; or use a hyphen to designate a range of port numbers. (Range: 1-28/52)

Show

- ◆ **MAC Address** – Physical address associated with the entry.
- ◆ **IP Address** – IP address corresponding to the client.
- ◆ **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 236: Configuring Static Bindings for IPv4 Source Guard

Security > IPv6 Source Guard > Static Binding

Action: Add ▼

Port: 1 ▼

VLAN: 1 ▼

MAC Address: 00-10-b5-f4-00-01 (xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx)

IPv6 Address: 2001::1


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 237: Displaying Static Bindings for IPv4 Source Guard



☐	VLAN	MAC Address	Interface	IP Address	Lease Time (sec)
☐	1	00-10-B5-F4-00-01	Unit 1 / Port 1	192.168.0.23	0

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** – DHCP-Snooping

Web Interface

To display the binding table for IP Source Guard:

1. Click Security, IP Source Guard, Dynamic Binding.
2. Mark the search criteria, and enter the required values.
3. Click Query

Figure 238: Showing the IPv4 Source Guard Binding Table

Security > IP Source Guard > Dynamic Binding

Query by:

☐ Port 1

☐ VLAN 1

☐ MAC Address (xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx)

☐ IP Address

Query

Dynamic Binding List Total: 3

VLAN	MAC Address	Interface	IP Address	Type
1	00-10-B5-F4-00-01	Unit 1 / Port 2	10.2.44.96	DHCP-Snooping
1	00-10-B5-F4-00-02	Unit 1 / Port 4	10.2.44.97	DHCP-Snooping
2	00-10-B5-F4-00-03	Unit 1 / Port 7	10.2.44.98	DHCP-Snooping

IPv6 Source Guard

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

Configuring Ports for IPv6 Source Guard

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

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

Command Usage

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

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

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier (Range: 1-28/52)
- ◆ **Filter Type** – Configures the switch to filter inbound traffic based on the following options. (Default: Disabled)
 - **Disabled** – Disables IPv6 source guard filtering on the port.
 - **SIP** – Enables traffic filtering based on IPv6 global unicast source IPv6 addresses stored in the binding table.
- ◆ **Max Binding Entry** – The maximum number of entries that can be bound to an interface. (ACL Table: 1-16, default: 5; MAC Table: 1-1024, default: 1024)
 - 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 381).

- 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 239: Setting the Filter Type for IPv6 Source Guard

The screenshot shows the 'Security > IPv6 Source Guard > Port Configuration' page. It features a 'Port Configuration List' with a total of 28 entries. The table has three columns: 'Port', 'Filter Type', and 'Max Binding Entry (1-5)'. There are five rows visible, each representing a port (1-5) with a 'Disabled' filter type and a '5' max binding entry. The interface includes a breadcrumb trail, a search icon, and pagination controls (1, 2, 3).

Port	Filter Type	Max Binding Entry (1-5)
1	Disabled	5
2	Disabled	5
3	Disabled	5
4	Disabled	5
5	Disabled	5

Configuring Static Bindings for IPv6 Source Guard

Use the Security > IPv6 Source Guard > Static Binding page to bind a static address to a port. Table entries include a MAC address, IPv6 global unicast address, entry type (Static-IPv6-SG-Binding, Dynamic-ND-Binding, Dynamic-DHCPv6-Binding), VLAN identifier, and port identifier.

Command Usage

- ◆ Traffic filtering is based only on the source IPv6 address, VLAN ID, and port number.

- ◆ Static addresses entered in the source guard binding table are automatically configured with an infinite lease time.
- ◆ When source guard is enabled, traffic is filtered based upon dynamic entries learned via ND snooping, DHCPv6 snooping, or static addresses configured in the source guard binding table.
- ◆ An entry with same MAC address and a different VLAN ID cannot be added to the binding table.
- ◆ Static bindings are processed as follows:
 - If there is no entry with same MAC address and IPv6 address, a new entry is added to binding table using static IPv6 source guard binding.
 - If there is an entry with same MAC address and IPv6 address, and the type of entry is static IPv6 source guard binding, then the new entry will replace the old one.
 - If there is an entry with same MAC address and IPv6 address, and the type of the entry is either a dynamic ND snooping binding or DHCPv6 snooping binding, then the new entry will replace the old one and the entry type will be changed to static IPv6 source guard binding.
 - Only unicast addresses are accepted for static bindings.

Parameters

These parameters are displayed:

Add

- ◆ **Port** – The port to which a static entry is bound.
- ◆ **VLAN** – ID of a configured VLAN (Range: 1-4094)
- ◆ **MAC Address** – A valid unicast MAC address.
- ◆ **IPv6 Address** – A valid global unicast IPv6 address. This address must be entered according to RFC 2373 “IPv6 Addressing Architecture,” using 8 colon-separated 16-bit hexadecimal values. One double colon may be used in the address to indicate the appropriate number of zeros required to fill the undefined fields.

Show

- ◆ **VLAN** – VLAN to which this entry is bound.
- ◆ **MAC Address** – Physical address associated with the entry.
- ◆ **Interface** – The port to which this entry is bound.
- ◆ **IPv6 Address** – IPv6 address corresponding to the client.

- ◆ **Type** – Shows the entry type:
 - **DHCP** – Dynamic DHCPv6 binding, stateful address.
 - **ND** – Dynamic Neighbor Discovery binding, stateless address.
 - **STA** – Static IPv6 Source Guard binding.

Web Interface

To configure static bindings for IPv6 Source Guard:

1. Click Security, IPv6 Source Guard, Static Configuration.
2. Select Add from the Action list.
3. Enter the required bindings for each port.
4. Click Apply

Figure 240: Configuring Static Bindings for IPv6 Source Guard

To display static bindings for IPv6 Source Guard:

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

Figure 241: Displaying Static Bindings for IPv6 Source Guard

VLAN	MAC Address	Interface	IPv6 Address	Type
1	00-10-B5-F4-00-01	Eth 1/2	10.2.44.96	STA
1	00-10-B5-F4-00-02	Eth 1/4	10.2.44.97	DHCP
2	00-10-B5-F4-00-03	Eth 1/7	10.2.44.98	ND

Displaying Information for Dynamic IPv6 Source Guard Bindings

Use the Security > IPv6 Source Guard > Dynamic Binding page to display the source-guard binding table for a selected interface.

Parameters

These parameters are displayed:

Query by

- ◆ **Port** – A port on this switch.
- ◆ **VLAN** – ID of a configured VLAN (Range: 1-4094)
- ◆ **MAC Address** – A valid unicast MAC address.
- ◆ **IPv6 Address** – A valid global unicast IPv6 address.

Dynamic Binding List


- ◆ **VLAN** – VLAN to which this entry is bound.
- ◆ **MAC Address** – Physical address associated with the entry.
- ◆ **Interface** – Port to which this entry is bound.
- ◆ **IPv6 Address** – IPv6 address corresponding to the client.
- ◆ **Type** – Shows the entry type:
 - **DHCP** – Dynamic DHCPv6 binding, stateful address.
 - **ND** – Dynamic Neighbor Discovery binding, stateless address.

Web Interface

To display the binding table for IPv6 Source Guard:

1. Click Security, IPv6 Source Guard, Dynamic Binding.
2. Mark the search criteria, and enter the required values.
3. Click Query

Figure 242: Showing the IPv6 Source Guard Binding Table



Security > IPv6 Source Guard > Dynamic Binding

Query by:

☐ Port

☐ VLAN

☐ MAC Address

☐ IPv6 Address

Query

Dynamic Binding List Total: 3

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

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 359](#)). 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 338](#)).

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.
 - **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 243: Configuring Global Settings for ARP Inspection

Security > ARP Inspection

Step: 1. Configure General

ARP Inspection Status ☒ Enabled

ARP Inspection Validation ☒ Dst-MAC ☒ IP ☒ Allow Zeros ☒ Src-MAC

Log Message Number (0-256) 20

Log Interval (0-86400) 10 sec

Apply Revert

Configuring VLAN Settings for ARP Inspection

Use the Security > ARP Inspection (Configure VLAN) page to enable ARP inspection for any VLAN and to specify the ARP ACL to use.

Command Usage

ARP Inspection VLAN Filters (ACLs)

- ◆ By default, no ARP Inspection ACLs are configured and the feature is disabled.
- ◆ ARP Inspection ACLs are configured within the ARP ACL configuration page (see [page 338](#)).
- ◆ 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** – Selects any configured VLAN. (Range: 1-4094; Default: 1)
- ◆ **DAI Status** – Enables Dynamic ARP Inspection for the selected VLAN. (Default: Disabled)
- ◆ **ACL Name** – Allows selection of any configured ARP ACLs. (Default: None)
- ◆ **ACL Status** – 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 244: Configuring VLAN Settings for ARP Inspection

Security > ARP Inspection

Step: 2. Configure VLAN

ARP Inspection VLAN List Total: 2

VLAN	DAI Status	ACL Name	ACL Status
1	☒ Enabled	☐ arp ▼	☐ Static
2	☒ Enabled	☐ arp ▼	☐ 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 the CPU per second on trusted or untrusted ports.
(Range: 0-2048; Default: 15)

Setting the rate limit to “0” means that there is no restriction on the number of ARP packets that can be processed by the CPU.

The switch will drop all ARP packets received on a port which exceeds the configured ARP-packets-per-second rate limit.

Web Interface

To configure interface settings for ARP Inspection:

1. Click Security, ARP Inspection.
2. Select Configure Interface from the Step list.
3. Specify any untrusted ports which require ARP inspection, and adjust the packet inspection rate.
4. Click Apply.

Figure 245: Configuring Interface Settings for ARP Inspection

Security > ARP Inspection

Step: 3. Configure Interface

Interface: ☒ Port ☐ Trunk

Port Configuration List Total: 28

Port	Trust Status	Packet Rate Limit (0-2048 pps)
1	☒ Enabled	☒ 15
2	☐ Enabled	☒ 15
3	☐ Enabled	☒ 15
4	☐ Enabled	☒ 15
5	☐ Enabled	☒ 15

Displaying ARP Inspection Statistics Use the Security > ARP Inspection (Show Information - Show Statistics) page to display statistics about the number of ARP packets processed, or dropped for various reasons.

Parameters

These parameters are displayed:

Table 24: ARP Inspection Statistics

Parameter	Description
Received ARP packets before ARP inspection rate limit	Count of ARP packets received but not exceeding the ARP Inspection rate limit.
Dropped ARP packets in the process of ARP inspection rate limit	Count of ARP packets exceeding (and dropped by) ARP rate limiting.
Total ARP packets processed by ARP inspection	Count of all ARP packets processed by the ARP Inspection engine.
ARP packets dropped by additional validation (Src-MAC)	Count of packets that failed the source MAC address test.
ARP packets dropped by additional validation (Dst-MAC)	Count of packets that failed the destination MAC address test.
ARP packets dropped by additional validation (IP)	Count of ARP packets that failed the IP address test.
ARP packets dropped by ARP ACLs	Count of ARP packets that failed validation against ARP ACL rules.
ARP packets dropped by DHCP snooping	Count of packets that failed validation against the DHCP Snooping Binding database.

Web Interface

To display statistics for ARP Inspection:

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

Figure 246: Displaying Statistics for ARP Inspection

Security > ARP Inspection	
Step:	4. Show Information
Action:	Show Statistics
Received ARP packets before ARP inspection rate limit	1000
Dropped ARP packets in processing ARP inspection rate limit	5
Total ARP packets processed by ARP inspection	200
ARP packets dropped by additional validation (Src-MAC)	300
ARP packets dropped by additional validation (Dst-MAC)	2000
ARP packets dropped by additional validation (IP)	100
ARP packets dropped by ARP ACLs	5
ARP packets dropped by DHCP snooping	5

Displaying the ARP Inspection Log

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

Parameters

These parameters are displayed:

Table 25: 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 247: Displaying the ARP Inspection Log

Security > ARP Inspection

Step: 4. Show Information Action: Show Log

ARP Inspection Log List Total: 2

VLAN ID	Port	Src. IP Address	Dst. IP Address	Src. MAC Address	Dst. MAC Address
1	15	192.168.1.1	192.168.1.5	11-22-33-44-55-66	AA-BB-CC-DD-EE-FF
1	17	192.168.1.3	192.168.1.23	11-4E-33-75-55-BB	A0-3B-C9-DD-4E-1F

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: 1-28/52)
- ◆ **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 248: Configuring Discarding or Forwarding of CDP/PVST Packets

Security > Application Filter		
Application Filter List Total: 28		
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 or PoE.
- ◆ **Ethernet Ring Protection Switching (ERPS)** – Configures a protection switching mechanism and protocol for Ethernet layer network rings.**Operation, Administration and Maintenance (OAM)** – Provides remote management tools required to monitor and maintain the links to subscriber CPEs (Customer Premise Equipment).
- ◆ **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.

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:

- ◆ **Status** – Enables/disables the logging of debug or error messages to the logging process. (Default: Enabled)
- ◆ **History Flash Level** – Limits log messages saved to the switch's permanent flash memory for all levels up to the specified level. For example, if level 3 is specified, all messages from level 0 to level 3 will be logged to flash. (Range: 0-7, Default: 3)

Table 26: Logging Levels

Level	Severity Name	Description
7	Debug	Debugging messages
6	Informational	Informational messages only
5	Notice	Normal but significant condition, such as cold start
4	Warning	Warning conditions (e.g., return false, unexpected return)
3	Error	Error conditions (e.g., invalid input, default used)
2	Critical	Critical conditions (e.g., memory allocation, or free memory error - resource exhausted)

* There are only Level 2, 5 and 6 error messages for the current firmware release.

Table 26: Logging Levels (Continued)

Level	Severity Name	Description
1	Alert	Immediate action needed
0	Emergency	System unusable

* There are only Level 2, 5 and 6 error messages for the current firmware release.

- ◆ **History 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 249: 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 250: 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 251: Configuring Settings for Remote Logging of Error Messages

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

- Remote Log Status:** A checkbox labeled 'Enabled' which is currently unchecked.
- Logging Facility:** A dropdown menu showing '23 - Local use 7'.
- Logging Trap Level:** A dropdown menu showing '0 - System unusable'.
- Server IP Address 1:** A text input field containing '192.168.0.4'.
- Server IP Address 2:** An empty text input field.
- Server IP Address 3:** An empty text input field.
- Server IP Address 4:** An empty text input field.
- Server IP Address 5:** An empty text input field.
- Port:** A column of text input fields. The first field (corresponding to IP 1) contains '514'. The others are empty.
- Buttons:** 'Apply' and 'Revert' buttons at the bottom right.

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 396](#)) 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. (Range: 1-41 characters)
- ◆ **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 669](#)), and one or more DNS servers specified (see “[Configuring a List of Name Servers](#)” on [page 672](#), or “[Configuring Static DNS Host to Address Entries](#)” on [page 673](#)).

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 252: Configuring SMTP Alert Messages

Link Layer Discovery Protocol

Link Layer Discovery Protocol (LLDP) is used to discover basic information about neighboring devices on the local broadcast domain. LLDP is a Layer 2 protocol that uses periodic broadcasts to advertise information about the sending device. Advertised information is represented in Type Length Value (TLV) format according to the IEEE 802.1AB standard, and can include details such as device identification, capabilities and configuration settings. LLDP also defines how to store and maintain information gathered about the neighboring network nodes it discovers.

Link Layer Discovery Protocol - Media Endpoint Discovery (LLDP-MED) is an extension of LLDP intended for managing endpoint devices such as Voice over IP phones and network switches. The LLDP-MED TLVs advertise information such as network policy, power, inventory, and device location details. LLDP and LLDP-MED information can be used by SNMP applications to simplify troubleshooting, enhance network management, and maintain an accurate network topology.

Setting LLDP Timing Attributes

Use the Administration > LLDP (Configure Global) page to set attributes for general functions such as globally enabling LLDP on the switch, setting the message ageout time, and setting the frequency for broadcasting general advertisements or reports about changes in the LLDP MIB.

Parameters

These parameters are displayed:

- ◆ **LLDP** – Enables LLDP globally on the switch. (Default: Enabled)

- ◆ **Transmission Interval** – Configures the periodic transmit interval for LLDP advertisements. (Range: 5-32768 seconds; Default: 30 seconds)

- ◆ **Hold Time Multiplier** – Configures the time-to-live (TTL) value sent in LLDP advertisements as shown in the formula below. (Range: 2-10; Default: 4)

The time-to-live tells the receiving LLDP agent how long to retain all information pertaining to the sending LLDP agent if it does not transmit updates in a timely manner. Note that if the local interface attached to a remote device is shut down or otherwise disabled, information about the remote device is purged immediately.

TTL in seconds is based on the following rule:

minimum value ((Transmission Interval * Holdtime Multiplier), or 65535)

Therefore, the default TTL is $4 * 30 = 120$ seconds.

- ◆ **Delay Interval** – Configures a delay between the successive transmission of advertisements initiated by a change in local LLDP MIB variables. (Range: 1-8192 seconds; Default: 2 seconds)

The transmit delay is used to prevent a series of successive LLDP transmissions during a short period of rapid changes in local LLDP MIB objects, and to increase the probability that multiple, rather than single changes, are reported in each transmission.

This attribute must comply with the rule:

$(4 * \text{Delay Interval}) \leq \text{Transmission Interval}$

- ◆ **Reinitialization Delay** – Configures the delay before attempting to re-initialize after LLDP ports are disabled or the link goes down. (Range: 1-10 seconds; Default: 2 seconds)

When LLDP is re-initialized on a port, all information in the remote systems LLDP MIB associated with this port is deleted.

- ◆ **Notification Interval** – Configures the allowed interval for sending SNMP notifications about LLDP MIB changes. (Range: 5-3600 seconds; Default: 5 seconds)

This parameter only applies to SNMP applications which use data stored in the LLDP MIB for network monitoring or management.

Information about changes in LLDP neighbors that occur between SNMP notifications is not transmitted. Only state changes that exist at the time of a notification are included in the transmission. An SNMP agent should therefore periodically check the value of `IldpStatsRemTableLastChangeTime` to detect any `IldpRemTablesChange` notification-events missed due to throttling or transmission loss.

- ◆ **MED Fast Start Count** – Configures the amount of LLDP MED Fast Start LLDPDUs to transmit during the activation process of the LLDP-MED Fast Start mechanism. (Range: 1-10 packets; Default: 4 packets)

The MED Fast Start Count parameter is part of the timer which ensures that the LLDP-MED Fast Start mechanism is active for the port. LLDP-MED Fast Start is

critical to the timely startup of LLDP, and therefore integral to the rapid availability of Emergency Call Service.

Web Interface

To configure LLDP timing attributes:

1. Click Administration, LLDP.
2. Select Configure Global from the Step list.
3. Enable LLDP, and modify any of the timing parameters as required.
4. Click Apply.

Figure 253: Configuring LLDP Timing Attributes

Administration > LLDP

Step: 1. Configure Global

LLDP ☒ Enabled

Transmission Interval (5-32768) sec

Hold Time Multiplier (2-10)

Delay Interval (1-8192) sec

Reinitialization Delay (1-10) sec

Notification Interval (5-3600) sec

MED Fast Start Count (1-10)

Note: The Transmission Interval must be greater than or equal to 4 times the Delay Interval.

Apply 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 in remote neighbors. (Default: Enabled)

This option sends out SNMP trap notifications to designated target stations at the interval specified by the Notification Interval in the preceding section. Trap notifications include information about state changes in the LLDP MIB (IEEE 802.1AB), the LLDP-MED MIB (ANSI/TIA-1057), or vendor-specific LLDP-EXT-DOT1 and LLDP-EXT-DOT3 MIBs.

For information on defining SNMP trap destinations, see [“Specifying Trap Managers” on page 446](#).

Information about additional changes in LLDP neighbors that occur between SNMP notifications is not transmitted. Only state changes that exist at the time of a trap notification are included in the transmission. An SNMP agent should therefore periodically check the value of `IldpStatsRemTableLastChangeTime` to detect any `IldpRemTablesChange` notification-events missed due to throttling or transmission loss.

- ◆ **MED Notification** – Enables the transmission of SNMP trap notifications about LLDP-MED changes. (Default: Disabled)
- ◆ **Basic Optional TLVs** – Configures basic information included in the TLV field of advertised messages.

- **Management Address** – The management address protocol packet includes the IPv4 address of the switch. If no management address is available, the address should be the MAC address for the CPU or for the port sending this advertisement. (Default: Enabled)

The management address TLV may also include information about the specific interface associated with this address, and an object identifier indicating the type of hardware component or protocol entity associated with this address. The interface number and OID are included to assist SNMP applications in the performance of network discovery by indicating enterprise specific or other starting points for the search, such as the Interface or Entity MIB.

Since there are typically a number of different addresses associated with a Layer 3 device, an individual LLDP PDU may contain more than one management address TLV.

Every management address TLV that reports an address that is accessible on a port and protocol VLAN through the particular port should be accompanied by a port and protocol VLAN TLV that indicates the VLAN identifier (VID) associated with the management address reported by this TLV.

- **Port Description** – The port description is taken from the `ifDescr` object in RFC 2863, which includes information about the manufacturer, the product name, and the version of the interface hardware/software. (Default: Enabled)
- **System Capabilities** – The system capabilities identifies the primary function(s) of the system and whether or not these primary functions are enabled. The information advertised by this TLV is described in IEEE 802.1AB. (Default: Enabled)
- **System Description** – The system description is taken from the `sysDescr` object in RFC 3418, which includes the full name and version identification

of the system's hardware type, software operating system, and networking software. (Default: Enabled)

- **System Name** – The system name is taken from the sysName object in RFC 3418, which contains the system's administratively assigned name. To configure the system name, see [“Displaying System Information” on page 72](#). (Default: Enabled)
- ◆ **802.1 Organizationally Specific TLVs** – Configures IEEE 802.1 information included in the TLV field of advertised messages.
 - **Protocol Identity** – The protocols that are accessible through this interface (see [“Protocol VLANs” on page 185](#)). (Default: Enabled)
 - **VLAN ID** – The port's default VLAN identifier (PVID) indicates the VLAN with which untagged or priority-tagged frames are associated (see [“IEEE 802.1Q VLANs” on page 159](#)). (Default: Enabled)
 - **VLAN Name** – The name of all VLANs to which this interface has been assigned (see [“IEEE 802.1Q VLANs” on page 159](#)). (Default: Enabled)
 - **Port and Protocol VLAN ID** – The port-based protocol VLANs configured on this interface (see [“Protocol VLANs” on page 185](#)). (Default: Enabled)
- ◆ **802.3 Organizationally Specific TLVs** – Configures IEEE 802.3 information included in the TLV field of advertised messages.
 - **Link Aggregation** – The link aggregation capabilities, aggregation status of the link, and the IEEE 802.3 aggregated port identifier if this interface is currently a link aggregation member. (Default: Enabled)
 - **Max Frame Size** – The maximum frame size. (See [“Configuring Support for Jumbo Frames” on page 74](#) for information on configuring the maximum frame size for this switch. (Default: Enabled)
 - **MAC/PHY Configuration/Status** – The MAC/PHY configuration and status which includes information about auto-negotiation support/capabilities, and operational Multistation Access Unit (MAU) type. (Default: Enabled)
 - **PoE** – Power-over-Ethernet capabilities, including whether or not PoE is supported, currently enabled, if the port pins through which power is delivered can be controlled, the port pins selected to deliver power, and the power class. (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** – 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 DHCP server.
 - Location of network element closest to client.
 - Location of client. (This is the default.)

Web Interface

To configure LLDP interface attributes:

1. Click Administration, LLDP.
2. Select Configure Interface from the Step list.
3. Select Configure General from the Action list.
4. Select an interface from the Port or Trunk list.
5. Set the LLDP transmit/receive mode, specify whether or not to send SNMP trap messages, and select the information to advertise in LLDP messages.
6. Click Apply.

Figure 254: Configuring LLDP Interface Attributes

Configuring LLDP Interface Civic-Address

Use the Administration > LLDP (Configure Interface – Add CA-Type) page to specify the physical location of the device attached to an interface.

Command Usage

- ◆ Use the Civic Address type (CA-Type) to advertise the physical location of the device attached to an interface, including items such as the city, street number, building and room information. The address location is specified as a type and value pair, with the civic address type defined in RFC 4776. The following table describes some of the CA type numbers and provides examples.

Table 27: LLDP MED Location CA Types

CA Type	Description	CA Value Example
1	National subdivisions (state, canton, province)	California
2	County, parish	Orange
3	City, township	Irvine
4	City division, borough, city district	West Irvine
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

Table 27: LLDP MED Location CA Types (Continued)

CA Type	Description	CA Value Example
21	Landmark or vanity address	Tech Center
26	Unit (apartment, suite)	Apt 519
27	Floor	5
28	Room	509B

- ◆ Any number of CA type and value pairs can be specified for the civic address location, as long as the total does not exceed 250 characters.

Parameters

These parameters are displayed:

- ◆ **CA-Type** – Descriptor of the data civic address value. (Range: 0-255)
- ◆ **CA-Value** – Description of a location. (Range: 1-32 characters)

Web Interface

To specify the physical location of the attached device:

1. Click Administration, LLDP.
2. Select Configure Interface from the Step list.
3. Select Add CA-Type from the Action list.
4. Select an interface from the Port or Trunk list.
5. Specify a CA-Type and CA-Value pair.
6. Click Apply.


Figure 255: Configuring the Civic Address for an LLDP Interface

The screenshot shows a web interface titled "Administration > LLDP". It has a "Step:" dropdown menu set to "2. Configure Interface" and an "Action:" dropdown menu set to "Add CA-Type". Below these, there are two radio buttons for "Interface": "Port" (selected) and "Trunk". Next to "Port" is a dropdown menu showing "1". Below the interface selection, there are two text input fields: "CA-Type (0-255)" with the value "1" and "CA-Value" with the value "California". At the bottom right, there are two buttons: "Apply" and "Revert".

To show the physical location of the attached device:

1. Click Administration, LLDP.
2. Select Configure Interface from the Step list.
3. Select Show CA-Type from the Action list.
4. Select an interface from the Port or Trunk list.

Figure 256: 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 28: Chassis ID Subtype

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

Table 28: Chassis ID Subtype (Continued)

ID Basis	Reference
Interface name	ifName (IETF RFC 2863)
Locally assigned	locally assigned

- ◆ **Chassis ID** – An octet string indicating the specific identifier for the particular chassis in this system.
- ◆ **System Name** – A string that indicates the system’s administratively assigned name (see [“Displaying System Information” on page 72](#)).
- ◆ **System Description** – A textual description of the network entity. This field is also displayed by the **show system** command.
- ◆ **System Capabilities Supported** – The capabilities that define the primary function(s) of the system.

Table 29: System Capabilities

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

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

Interface Settings

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

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

Interface Details

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

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

Table 30: Port ID Subtype

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

- ◆ **Port/Trunk ID** – A string that contains the specific identifier for the local interface based on interface subtype used by this switch.
- ◆ **Port/Trunk Description** – A string that indicates the port or trunk description. If RFC 2863 is implemented, the ifDescr object should be used for this field.
- ◆ **MED Capability** – The supported set of capabilities that define the primary function(s) of the interface:
 - LLDP-MED Capabilities
 - Network Policy
 - Location Identification
 - Extended Power via MDI – PSE
 - Extended Power via MDI – PD
 - Inventory

Web Interface

To display LLDP information for the local device:

1. Click Administration, LLDP.
2. Select Show Local Device Information from the Step list.

3. Select General, Port, Port Details, Trunk, or Trunk Details.

Figure 257: Displaying Local Device Information for LLDP (General)

Administration > LLDP

Step: 3. Show Local Device Information

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

LLDP Local Device Information

Chassis Type	MAC Address
Chassis ID	70-72-CF-83-34-66
System Name	
System Description	ECS4620-S2T
System Capabilities Supported	Bridge, Router
System Capabilities Enabled	Bridge, Router
Management Address	192.168.0.2 (IPv4)

Figure 258: Displaying Local Device Information for LLDP (Port)

Administration > LLDP

Step: 3. Show Local Device Information

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

LLDP Local Device Port List Total: 28

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

Figure 259: Displaying Local Device Information for LLDP (Port Details)

Administration > LLDP

Step: 3. Show Local Device Information

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

Port 1

LLDP Local Port Information Details

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

Displaying LLDP Remote Device Information

Use the Administration > LLDP (Show Remote Device Information) page to display information about devices connected directly to the switch's ports which are advertising information through LLDP, or to display detailed information about an LLDP-enabled device connected to a specific port on the local switch.

Parameters

These parameters are displayed:

Port

- ◆ **Local Port** – The local port to which a remote LLDP-capable device is attached.
- ◆ **Chassis ID** – An octet string indicating the specific identifier for the particular chassis in this system.
- ◆ **Port ID** – A string that contains the specific identifier for the port from which this LLDPDU was transmitted.
- ◆ **System Name** – A string that indicates the system's administratively assigned name.

Port Details

- ◆ **Port** – Port identifier on local switch.
- ◆ **Remote Index** – Index of remote device attached to this port.
- ◆ **Local Port** – The local port to which a remote LLDP-capable device is attached.
- ◆ **Chassis Type** – Identifies the chassis containing the IEEE 802 LAN entity associated with the transmitting LLDP agent. There are several ways in which a chassis may be identified and a chassis ID subtype is used to indicate the type of component being referenced by the chassis ID field. (See [Table 28, "Chassis ID Subtype," on page 409.](#))
- ◆ **Chassis ID** – An octet string indicating the specific identifier for the particular chassis in this system.
- ◆ **System Name** – A string that indicates the system's assigned name.
- ◆ **System Description** – A textual description of the network entity.
- ◆ **Port Type** – Indicates the basis for the identifier that is listed in the Port ID field. See [Table 30, "Port ID Subtype," on page 411.](#)
- ◆ **Port Description** – A string that indicates the port's description. If RFC 2863 is implemented, the ifDescr object should be used for this field.
- ◆ **Port ID** – A string that contains the specific identifier for the port from which this LLDPDU was transmitted.

- ◆ **System Capabilities Supported** – The capabilities that define the primary function(s) of the system. (See [Table 29, "System Capabilities," on page 410.](#))
- ◆ **System Capabilities Enabled** – The primary function(s) of the system which are currently enabled. (See [Table 29, "System Capabilities," on page 410.](#))
- ◆ **Management Address List** – The management addresses for this device. Since there are typically a number of different addresses associated with a Layer 3 device, an individual LLDP PDU may contain more than one management address TLV.

If no management address is available, the address should be the MAC address for the CPU or for the port sending this advertisement.

Port Details – 802.1 Extension Information

- ◆ **Remote Port VID** – The port's default VLAN identifier (PVID) indicates the VLAN with which untagged or priority-tagged frames are associated.
- ◆ **Remote Port-Protocol VLAN List** – The port-based protocol VLANs configured on this interface, whether the given port (associated with the remote system) supports port-based protocol VLANs, and whether the port-based protocol VLANs are enabled on the given port associated with the remote system.
- ◆ **Remote VLAN Name List** – VLAN names associated with a port.
- ◆ **Remote Protocol Identity List** – Information about particular protocols that are accessible through a port. This object represents an arbitrary local integer value used by this agent to identify a particular protocol identity, and an octet string used to identify the protocols associated with a port of the remote system.

Port Details – 802.3 Extension Port Information

- ◆ **Remote Port Auto-Neg Supported** – Shows whether the given port (associated with remote system) supports auto-negotiation.
- ◆ **Remote Port Auto-Neg Adv-Capability** – The value (bitmap) of the ifMauAutoNegCapAdvertisedBits object (defined in IETF RFC 3636) which is associated with a port on the remote system.

Table 31: Remote Port Auto-Negotiation Advertised Capability

Bit	Capability
0	other or unknown
1	10BASE-T half duplex mode
2	10BASE-T full duplex mode
3	100BASE-T4
4	100BASE-TX half duplex mode

Table 31: 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 ¹⁴

- ◆ **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.

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

Port Details – Network Policy¹⁴

- ◆ **Application Type** – The primary applications 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 Identification¹⁴

- ◆ **Location Data Format** – Any of these location ID data formats:
 - Coordinate-based LCI¹⁵ – Defined in RFC 3825, includes latitude resolution, latitude, longitude resolution, longitude, altitude type, altitude resolution, altitude, and datum.
 - Civic Address LCI¹⁵ – 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.”](#)

15. 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 – Inventory¹⁴

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

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

Top Screenshot: The interface shows the breadcrumb "Administration > LLDP". The "Step:" dropdown is set to "4. Show Remote Device Information". Below this, there are four radio buttons: "Port" (selected), "Port Details", "Trunk", and "Trunk Details". A table titled "LLDP Remote Device Port List" with "Total: 1" is displayed. The table has four columns: "Local Port", "Chassis ID", "Port ID", and "System Name". The single row shows "1" for Local Port, "70-72-CF-32-DD-FD" for Chassis ID, "70-72-CF-32-DD-FF" for Port ID, and "RD93" for System Name.

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

Bottom Screenshot: The interface shows the same breadcrumb "Administration > LLDP". The "Step:" dropdown is set to "4. Show Remote Device Information". Below this, there are four radio buttons: "Port", "Port Details" (selected), "Trunk", and "Trunk Details". There are two dropdown menus: "Port" with "1" selected and "Remote Index" with "2" selected. A "Query" button is located at the bottom right.

Figure 261: Displaying Remote Device Information for LLDP (Port Details)

Administration > LLDP

Step: 4. Show Remote Device Information

Port
23
Remote Index
4
Query

LLDP Remote Device Port Information

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

Management Address List Total: 1

Address	Address Type
192.168.0.3	IPv4 Address

802.1 Extension Information

Remote Port VID	1

802.1 Extension Information

Remote Port VID	1

Remote Port-Protocol VLAN List Total: 1

VLAN	Support	Status
1	Yes	Enabled

Remote VLAN Name List Total: 1

VLAN	Name
1	DefaultVlan

Remote Protocol Identity List Total: 1

Remote Protocol Identity (Hex)
88-CC

802.3 Extension Port Information

Remote Port Auto-Neg Supported	Yes	Remote Port Auto-Neg Status	Enabled
Remote Port Auto-Neg Adv-Capability	0000	Remote Port MAU Type	6

802.3 Extension Power Information

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

802.3 Extension Trunk Information

Remote Link Aggregation Capable	Yes	Remote Link Aggregation Status	Disabled
Remote Link Port ID	0		

802.3 Extension Frame Information

Remote Max Frame Size	1518

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 262: Displaying LLDP Device Statistics (General)

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

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 263: Displaying LLDP Device Statistics (Port)

The screenshot shows the 'Administration > LLDP' web interface with the 'Port' tab selected. The 'Step' dropdown is set to '5. Show Device Statistics'. The 'Port' dropdown is set to '3'. The 'LLDP Device Port Statistics' section displays the following data:

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

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

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, power is not supplied.

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 during bootup.

Setting the Switch's Overall PoE Power Budget

Use the Administration > PoE > PSE (Configure Global) page to set 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: 370 Watts)
- ◆ **PoE Maximum Allocation 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. (Range: 37000 - 370000 milliwatts; Default: 370 Watts)
- ◆ **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 set the overall PoE power budget for switch:

1. Click Administration, PoE.
2. Select Configure Global from the Step list.
3. Set the maximum PoE power provided by the switch.
4. Click Apply.

Figure 264: Setting the Switch's PoE Budget

The screenshot shows a web interface for configuring PoE settings. The breadcrumb navigation at the top reads "Administration > PoE > PSE". Below this, a "Step:" dropdown menu is set to "1. Configure Global". The main configuration area contains several fields: "PoE Maximum Available Power" is set to "370000 milliwatts"; "PoE Maximum Allocation Power (37000-370000)" has a text input field containing "370000" followed by "milliwatts"; "System Operation Status" is set to "On"; "PoE Power Consumption" is set to "0 milliwatts"; and "Software Version" is set to "Version 1.7". At the bottom right of the configuration area are two buttons: "Apply" and "Revert".

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

- ◆ This switch supports both the IEEE 802.3af PoE and IEEE 802.3at-2009 PoE Plus standards. To ensure that the correct power is supplied to powered devices (PD) compliant with these standards, the first detection pulse from the switch is based on 802.3af to which the 802.3af PDs will respond normally. It then sends a second PoE Plus pulse that causes an 802.3at PD to respond as a Class 4 device and draw Class 4 current. Afterwards, the switch exchanges information with the PD such as duty-cycle, peak and average power needs.
- ◆ The total PoE power delivered by all ports cannot exceed the maximum power budget of 370W. All the RJ-45 ports support both the IEEE 802.3af and IEEE 802.3at standards. The maximum number of ports which can supply power simultaneously at the specified levels are shown in the following table.

Table 32: Maximum Number of Ports Providing Simultaneous Power

Switch	34.2W (802.3at)	15.4W (802.3af)	7.5W (802.3af)
ECS4120-28P	10	24	24

- ◆ 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 or to the overall switch, 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. For example:
 - If a device is connected to a low-priority port and causes the switch to exceed its budget, power to this port is not turned on.
 - If a device is connected to a critical or high-priority port and 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.
 - If a device is connected to a port after the switch has finished booting up and would cause the switch to exceed its budget, power will not be provided to that port regardless of its priority setting.
 - If priority is not set for any ports, and PoE consumption exceeds the maximum power provided by the switch, power is shut down starting from the higher-numbered ports.

Parameters

These parameters are displayed:

- ◆ **Port** – The port number on the switch. (Range: 1-28/52)
- ◆ **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.
- ◆ **Priority** – Sets the power priority for a port. (Options: Low, High or Critical; Default: Low)
- ◆ **Power Allocation** – Sets the power budget for a port. (Range: 3400-34200 milliwatts; Default: 34200 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 265: Setting a Port's PoE Budget

The screenshot shows a web interface titled "Administration > PoE > PSE". Below the title is a "Step:" dropdown menu set to "2. Configure Interface". The main section is titled "PoE Port List" with a "Total: 24" indicator. It contains a table with 10 rows, each representing a port. The columns are: Port, Admin Status, Mode, Time-Range Name, Time-Range Status, Priority, Power Allocation, and Power Consumption (milliwatts). All ports are currently "Enabled", "off" mode, "R&D" time-range, "NA" status, "Low" priority, and have a "Power Allocation" of 34200 and "Power Consumption" of 0. At the bottom of the table are "Apply" and "Revert" buttons.

Port	Admin Status	Mode	Time-Range Name	Time-Range Status	Priority	Power Allocation	Power Consumption (milliwatts)
1	☒ Enabled	off	☐ R&D ▼	NA	Low ▼	34200	0
2	☒ Enabled	off	☐ R&D ▼	NA	Low ▼	34200	0
3	☒ Enabled	off	☐ R&D ▼	NA	Low ▼	34200	0
4	☒ Enabled	off	☐ R&D ▼	NA	Low ▼	34200	0
5	☒ Enabled	off	☐ R&D ▼	NA	Low ▼	34200	0
6	☒ Enabled	off	☐ R&D ▼	NA	Low ▼	34200	0
7	☒ Enabled	off	☐ R&D ▼	NA	Low ▼	34200	0
8	☒ Enabled	off	☐ R&D ▼	NA	Low ▼	34200	0
9	☒ Enabled	off	☐ R&D ▼	NA	Low ▼	34200	0
10	☒ Enabled	off	☐ R&D ▼	NA	Low ▼	34200	0

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 33: SNMPv3 Security Models and Levels

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



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

Command Usage

Configuring SNMPv1/2c Management Access

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

1. Use the Administration > SNMP (Configure Global) page to enable SNMP on the switch, and to enable trap messages.
2. Use the Administration > SNMP (Configure User - Add Community) page to configure the community strings authorized for management access.
3. Use the Administration > SNMP (Configure Trap) page to specify trap managers so that key events are reported by this switch to your management station.

Configuring SNMPv3 Management Access

1. Use the Administration > SNMP (Configure Global) page to enable SNMP on the switch, and to enable trap messages.
2. Use the Administration > SNMP (Configure Trap) page to specify trap managers so that key events are reported by this switch to your management station.
3. Use the Administration > SNMP (Configure Engine) page to change the local engine ID. If you want to change the default engine ID, it must be changed before configuring other parameters.
4. Use the Administration > SNMP (Configure View) page to specify read and write access views for the switch MIB tree.
5. Use the Administration > SNMP (Configure User) page to configure SNMP user groups with the required security model (i.e., SNMP v1, v2c or v3) and security level (i.e., authentication and privacy).
6. Use the Administration > SNMP (Configure Group) page to assign SNMP users to groups, along with their specific authentication and privacy passwords.

Configuring Global Settings for SNMP

Use the Administration > SNMP (Configure Global) page to enable SNMPv3 service for all management clients (i.e., versions 1, 2c, 3), and to enable trap messages.

Parameters

These parameters are displayed:

- ◆ **Agent Status** – Enables SNMP on the switch. (Default: Enabled)

- ◆ **Authentication Traps**¹⁶ – 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 266: Configuring Global Settings for SNMP

The screenshot shows a web interface for configuring SNMP. At the top, it says 'Administration > SNMP'. Below that is a 'Step:' dropdown menu set to '1. Configure Global'. The main area contains two settings: 'Agent Status' and 'Authentication Traps', both with checked checkboxes and the word 'Enabled' next to them. At the bottom right, there are 'Apply' and 'Revert' buttons.

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)
- ◆ **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.

16. 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 433).

Web Interface

To set a community access string:

1. Click Administration, SNMP.
2. Select Configure User from the Step list.
3. Select Add Community from the Action list.
4. Add new community strings as required, and select the corresponding access rights from the Access Mode list.
5. Click Apply

Figure 267: Setting Community Access Strings

The screenshot shows the 'Administration > SNMP' page. The 'Step' dropdown is set to '2. Configure Community' and the 'Action' dropdown is set to 'Add'. Below these, there is a 'Community String' input field containing 'spidermain' and an 'Access Mode' 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 User from the Step list.
3. Select Show Community from the Action list.

Figure 268: Showing Community Access Strings

The screenshot shows the 'Administration > SNMP' page with the 'Action' dropdown set to 'Show'. Below the dropdowns, there is a table titled 'SNMP Community String List' with a 'Total: 3' indicator. The table has three columns: a checkbox column, 'Community String', and 'Access Mode'. There are three rows of data: 'public' (Read-Only), 'private' (Read/Write), and 'spidermain' (Read/Write). At the bottom right, there are 'Delete' and 'Revert' buttons.

	Community String	Access Mode
☐	public	Read-Only
☐	private	Read/Write
☐	spidermain	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 269: Configuring the Local Engine ID for SNMP

The screenshot shows the 'Administration > SNMP' web interface. At the top, there is a breadcrumb trail 'Administration > SNMP'. Below it, there are two dropdown menus: 'Step:' with '3. Configure Engine' selected, and 'Action:' with 'Set Engine ID' selected. The main configuration area contains two fields: 'Engine ID' with the value '800001030300e00c0000fa0000' and 'Engine Boots' with the value '8'. At the bottom right, there are two buttons: 'Default' and 'Save'.

Specifying a Remote Engine ID

Use the Administration > SNMP (Configure Engine - Add Remote Engine) page to configure a engine ID for a remote management station. To allow management access from an SNMPv3 user on a remote device, you must first specify the engine identifier for the SNMP agent on the remote device where the user resides. The remote engine ID is used to compute the security digest for authentication and encryption of packets passed between the switch and a user on the remote host.

Command Usage

- ◆ SNMP passwords are localized using the engine ID of the authoritative agent. For informs, the authoritative SNMP agent is the remote agent. You therefore need to configure the remote agent's SNMP engine ID before you can send proxy requests or informs to it. (See ["Configuring Remote SNMPv3 Users" on page 443.](#))

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 270: 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 "3. Configure Engine" selected, and "Action:" with "Add Remote Engine" selected. Below these, there are two input fields: "Remote Engine ID" with the value "543210000" and "Remote IP Host" with the value "192.168.1.19". At the bottom right, there are two buttons: "Apply" and "Revert".

To show the remote SNMP engine IDs:

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

Figure 271: Showing Remote Engine IDs for SNMP

Administration > SNMP

Step: 3. Configure Engine Action: Show Remote Engine

SNMPv3 Remote Engine List Total: 1

	Remote Engine ID	Remote IP Host
☐	5432100000	192.168.1.19

Delete Revert

Setting SNMPv3 Views Use the Administration > SNMP (Configure View) page to configure SNMPv3 views which are used to restrict user access to specified portions of the MIB tree. The predefined view "defaultview" includes access to the entire MIB tree.

Parameters

These parameters are displayed:

Add View

- ◆ **View Name** – The name of the SNMP view. A maximum of 32 views can be configured. (Range: 1-32 characters)
- ◆ **OID Subtree** – Specifies the initial object identifier of a branch within the MIB tree. Wild cards can be used to mask a specific portion of the OID string. Use the Add OID Subtree page to configure additional object identifiers. (Range: 1-64 characters)
- ◆ **Type** – Indicates if the object identifier of a branch within the MIB tree is included or excluded from the SNMP view.

Add OID Subtree

- ◆ **View Name** – Lists the SNMP views configured in the Add View page. A maximum of 32 views can be configured. (Range: 1-32 characters)
- ◆ **OID Subtree** – Adds an additional object identifier of a branch within the MIB tree to the selected View. Wild cards can be used to mask a specific portion of the OID string. (Range: 1-64 characters)
- ◆ **Type** – Indicates if the object identifier of a branch within the MIB tree is included or excluded from the SNMP view.

Web Interface

To configure an SNMP view of the switch's MIB database:

1. Click Administration, SNMP.
2. Select Configure View from the Step list.
3. Select Add View from the Action list.
4. Enter a view name and specify the initial OID subtree in the switch's MIB database to be included or excluded in the view. Use the Add OID Subtree page to add additional object identifier branches to the view.
5. Click Apply

Figure 272: Creating an SNMP View

Administration > SNMP

Step: 4. Configure View Action: Add View

View Name: ifEntry.a

OID Subtree: 1.3.6.1.2.1.2.2.1.1.*

Type: Included

Apply Revert

To show the SNMP views of the switch's MIB database:

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

Figure 273: Showing SNMP Views

Administration > SNMP

Step: 4. Configure View Action: Show View

SNMPv3 View List Total: 2

	View Name
☐	ifEntry.a
☐	defaultview

Delete Revert

To add an object identifier to an existing SNMP view of the switch's MIB database:

1. Click Administration, SNMP.
2. Select Configure View from the Step list.

3. Select Add OID Subtree from the Action list.
4. Select a view name from the list of existing views, and specify an additional OID subtree in the switch's MIB database to be included or excluded in the view.
5. Click Apply

Figure 274: Adding an OID Subtree to an SNMP View

The screenshot shows the 'Administration > SNMP' configuration page. At the top, there's a breadcrumb 'Administration > SNMP'. Below it, a 'Step' dropdown is set to '4. Configure View' and an 'Action' dropdown is set to 'Add OID Subtree'. The main form has three fields: 'View Name' with a dropdown showing 'ifEntry.a', 'OID Subtree' with a text input containing '1.3.6.1.2.1.2.2.1.2.*', and 'Type' with a dropdown showing 'Included'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 275: Showing the OID Subtree Configured for SNMP Views

The screenshot shows the 'Administration > SNMP' configuration page with the 'Action' dropdown set to 'Show OID Subtree'. The 'View Name' dropdown still shows 'ifEntry.a'. Below the form, there's a section titled 'SNMPv3 View OID Subtree List' with a 'Total: 2' indicator. It contains a table with two rows of configured OID subtrees. At the bottom right, there are 'Delete' and 'Revert' buttons.

	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

Configuring SNMPv3 Groups Use the Administration > SNMP (Configure Group) page to add an SNMPv3 group which can be used to set the access policy for its assigned users, restricting them to specific read, write, and notify views. You can use the pre-defined default groups or create new groups to map a set of SNMP users to SNMP views.

Parameters

These parameters are displayed:

- ◆ **Group Name** – The name of the SNMP group to which the user is assigned. A maximum of 22 groups can be configured. (Range: 1-32 characters)
- ◆ **Security Model** – The user security model; SNMP v1, v2c or v3.
- ◆ **Security Level** – The following security levels are only used for the groups assigned to the SNMP security model:
 - **noAuthNoPriv** – There is no authentication or encryption used in SNMP communications. (This is the default security level.)
 - **AuthNoPriv** – SNMP communications use authentication, but the data is not encrypted.
 - **AuthPriv** – SNMP communications use both authentication and encryption.
- ◆ **Read View** – The configured view for read access. (Range: 1-32 characters)
- ◆ **Write View** – The configured view for write access. (Range: 1-32 characters)
- ◆ **Notify View** – The configured view for notifications. (Range: 1-32 characters)

Table 34: Supported Notification Messages

Model	Level	Group
<i>RFC 1493 Traps</i>		
newRoot	1.3.6.1.2.1.17.0.1	The newRoot trap indicates that the sending agent has become the new root of the Spanning Tree; the trap is sent by a bridge soon after its election as the new root, e.g., upon expiration of the Topology Change Timer immediately subsequent to its election.
topologyChange	1.3.6.1.2.1.17.0.2	A topologyChange trap is sent by a bridge when any of its configured ports transitions from the Learning state to the Forwarding state, or from the Forwarding state to the Discarding state. The trap is not sent if a newRoot trap is sent for the same transition.

Table 34: Supported Notification Messages (Continued)

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.
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	1.3.6.1.4.1.259.6.10.120.2.1.0.1	This trap is sent when the power state changes.
swFanFailureTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.17	This trap is sent when the fan is failure.
swFanRecoverTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.18	This trap is sent when fan failure has recovered.
swPortSecurityTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.36	This trap is sent when the port is being intruded. This trap will only be sent when the portSecActionTrap is enabled.
swIpFilterRejectTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.40	This trap is sent when an incorrect IP address is rejected by the IP Filter.
swSmtppConnFailureTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.41	This trap is triggered if the SMTP system cannot open a connection to the mail server successfully.
swThermalRisingNotification	1.3.6.1.4.1.259.6.10.120.2.1.0.58	This trap is sent when the temperature is over the switchThermalActionRisingThreshold.
swThermalFallingNotification	1.3.6.1.4.1.259.6.10.120.2.1.0.59	This trap is sent when the temperature is below the switchThermalActionFallingThreshold.

Table 34: Supported Notification Messages (Continued)

Model	Level	Group
pethPsePortOnOffNotification	1.3.6.1.4.1.259.6.10.120.2.1.0.43	This notification indicates if a PSE port is delivering power to the PD. This notification should be sent on every status change except in searching mode.
pethPsePortPowerMaintenanceStatusNotification	1.3.6.1.4.1.259.6.10.120.2.1.0.44	This notification indicates a Port Change Status and should be sent on every status change.
pethMainPowerUsageOnNotification	1.3.6.1.4.1.259.6.10.120.2.1.0.45	This notification indicates PSE threshold usage indication is on; and the power usage is above the threshold.
pethMainPowerUsageOffNotification	1.3.6.1.4.1.259.6.10.120.2.1.0.46	This notification indicates that the PSE threshold usage indication is off; and the usage power is below the threshold.
swAtcBcastStormAlarmFireTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.70	When broadcast traffic is detected as a storm, this trap is fired.
swAtcBcastStormAlarmClearTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.71	When a broadcast storm is detected as normal traffic, this trap is fired.
swAtcBcastStormTcApplyTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.72	When ATC is activated, this trap is fired.
swAtcBcastStormTcReleaseTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.73	When ATC is released, this trap is fired.
swAtcMcastStormAlarmFireTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.74	When multicast traffic is detected as the storm, this trap is fired.
swAtcMcastStormAlarmClearTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.75	When multicast storm is detected as normal traffic, this trap is fired.
swAtcMcastStormTcApplyTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.76	When ATC is activated, this trap is fired.
swAtcMcastStormTcReleaseTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.77	When ATC is released, this trap is fired.
stpBpduGuardPortShutdownTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.91	This trap will be sent when an interface is shut down because of BPDU guard.
swLoopbackDetectionTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.95	This trap is sent when loopback BPDUs have been detected.
networkAccessPortLinkDetectionTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.96	This trap is sent when a networkAccessPortLinkDetection event is triggered.
dot1agCfmMepUpTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.97	This trap is sent when a new remote MEP is discovered.
dot1agCfmMepDownTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.98	This trap is sent when port status or interface status TLV received from remote MEP indicates it is not up.
dot1agCfmConfigFailTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.99	This trap is sent when a MEP receives a CCM with MPID which already exists on the same MA in this switch.
dot1agCfmLoopFindTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.100	This trap is sent when a MEP receives its own CCMs.
dot1agCfmMepUnknownTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.101	This trap is sent when a CCM is received from an unexpected MEP.
dot1agCfmMepMissingTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.102	This trap is sent when the cross-check enable timer expires and no CCMs were received from an expected (configured) MEP.
dot1agCfmMaUpTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.103	This trap is sent when all expected remote MEPs are up.
autoUpgradeTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.104	This trap is sent when auto upgrade is executed.

Table 34: Supported Notification Messages (Continued)

Model	Level	Group
swCpuUtiRisingNotification	1.3.6.1.4.1.259.6.10.120.2.1.0.107	This notification indicates that the CPU utilization has risen from cpuUtiFallingThreshold to cpuUtiRisingThreshold.
swCpuUtiFallingNotification	1.3.6.1.4.1.259.6.10.120.2.1.0.108	This notification indicates that the CPU utilization has fallen from cpuUtiRisingThreshold to cpuUtiFallingThreshold.
swMemoryUtiRisingThreshold Notification	1.3.6.1.4.1.259.6.10.120.2.1.0.109	This notification indicates that the memory utilization has risen from memoryUtiFallingThreshold to memoryUtiRisingThreshold.
swMemoryUtiFallingThreshold Notification	1.3.6.1.4.1.259.6.10.120.2.1.0.110	This notification indicates that the memory utilization has fallen from memoryUtiRisingThreshold to memoryUtiFallingThreshold.
dhcpRogueServerAttackTrap	1.3.6.1.4.1.259.10.1.24.2.1.0.114	This trap is sent when receiving a DHCP packet from a rouge server.
macNotificationTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.138	This trap is sent when there are changes of the dynamic MAC addresses on the switch.
lbdDetectionTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.141	This trap is sent when a loopback condition is detected by LBD.
lbdRecoveryTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.142	This trap is sent when a recovery is done by LBD.
sfpThresholdAlarmWarnTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.189	This trap is sent when the SFP's A/D quantity is not within alarm/warning thresholds.
udldPortShutdownTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.192	This trap is sent when the port is shut down by UDLD.
userAuthenticationFailureTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.199	This trap will be triggered if authentication fails.
userAuthenticationSuccessTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.200	This trap will be triggered if authentication is successful.
loginTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.201	This trap is sent when user logs in.
logoutTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.202	This trap is sent when user logs out.
fileCopyTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.208	This trap is sent when file copy is executed. If the copy action is triggered by the system, the login user information (trapVarLoginUserName/ trapVarSessionType/trapVarLoginInetAddressTypes/ trapVarLoginInetAddress) will be a null value.
userauthCreateUserTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.209	This trap is sent when a user account is created.
userauthDeleteUserTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.210	This trap is sent when a user account is deleted.
userauthModifyUserPrivilegeTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.211	This trap is sent when user privilege is modified.
cpuGuardControlTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.213	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	1.3.6.1.4.1.259.6.10.120.2.1.0.214	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 276: Creating an SNMP Group

Administration > SNMP

Step: 5. Configure Group Action: Add

Group Name: secure-users

Security Model: v3

Security Level: authPriv

Read View: ☒ ifEntry.a ☐

Write View: ☒ ifEntry.a ☐

Notify View: ☒ ifEntry.a ☐

Apply Revert

To show SNMP groups:

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

Figure 277: Showing SNMP Groups

Administration > SNMP

Step: 5. Configure Group Action: Show

SNMPv3 Group List Total: 7

	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
☐	spidermain	v1	noAuthNoPriv	defaultview	defaultview	No notifyview specified
☐	spidermain	v2c	noAuthNoPriv	defaultview	defaultview	No notifyview specified
☒	secure-users	v3	authPriv	ifEntry.a	ifEntry.a	ifEntry.a

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. A maximum of three local users can be configured. (Range: 1-32 characters)
- ◆ **Group Name** – The name of the SNMP group to which the user is assigned. (Range: 1-32 characters)
- ◆ **Security Model** – The user security model; SNMP v1, v2c or v3.
- ◆ **Security Level** – The following security levels are only used for the groups assigned to the SNMP security model:
 - **noAuthNoPriv** – There is no authentication or encryption used in SNMP communications. (This is the default security level.)
 - **AuthNoPriv** – SNMP communications use authentication, but the data is not encrypted.
 - **AuthPriv** – SNMP communications use both authentication and encryption.
- ◆ **Authentication Protocol** – The method used for user authentication. (Options: MD5, SHA; Default: MD5)
- ◆ **Authentication Password** – A minimum of eight plain text characters is required. (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 278: Configuring Local SNMPv3 Users

The screenshot shows the 'Administration > SNMP' configuration page. The 'Step' dropdown is set to '6. Configure User' and the 'Action' dropdown is set to 'Add SNMPv3 Local User'. The form is divided into three sections: 'SNMPv3 User', 'User Authentication', and 'Data Privacy'. In the 'SNMPv3 User' section, 'User Name' is 'chris', 'Group Name' is 'public' (selected from a dropdown), 'Security Model' is 'v3', and 'Security Level' is 'authPriv'. In the 'User Authentication' section, 'Authentication Protocol' is 'MD5' and 'Authentication Password' is 'greenpeace'. In the 'Data Privacy' section, 'Privacy Protocol' is 'DES56' and 'Privacy Password' is 'einstien'. At the bottom right are 'Apply' and 'Revert' buttons.

To show local SNMPv3 users:

1. Click Administration, SNMP.
2. Select Configure User from the Step list.
3. Select Show SNMPv3 Local User from the Action list.

Figure 279: Showing Local SNMPv3 Users

The screenshot shows the 'Administration > SNMP' configuration page with the 'Action' dropdown set to 'Show SNMPv3 Local User'. Below the form is a table titled 'SNMPv3 Local User List' with a 'Total: 1' indicator. The table has columns for 'User Name', 'Group Name', 'Model', 'Level', 'Authentication', and 'Privacy'. There is one row for the user 'chris' in the 'public' group, using 'v3' model, 'authPriv' level, 'MD5' authentication, and 'DES56' privacy. At the bottom right are 'Delete' and 'Revert' buttons.

	User Name	Group Name	Model	Level	Authentication	Privacy
☐	chris	public	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 280: Changing a Local SNMPv3 User Group

The screenshot shows a web interface for configuring SNMPv3 users. The breadcrumb is 'Administration > SNMP'. The 'Step' is '6. Configure User' and the 'Action' is 'Change SNMPv3 Local User Group'. Under 'User Name', a dropdown menu shows 'chris'. Under 'Group Name', there are two radio buttons: 'bert' (which is selected) and 'public' (which is shown in a dropdown menu below it). 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 446](#) and [“Specifying a Remote Engine ID” on page 432](#).)

Parameters

These parameters are displayed:

- ◆ **User Name** – The name of user connecting to the SNMP agent.
(Range: 1-32 characters)
- ◆ **Group Name** – The name of the SNMP group to which the user is assigned.
(Range: 1-32 characters)
- ◆ **Remote IP** – The IPv4 address of the host (the targeted recipient).
- ◆ **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 281: Configuring Remote SNMPv3 Users

Administration > SNMP

Step: 6. Configure User Action: Add SNMPv3 Remote User

SNMPv3 User

User Name: mark

Group Name: ☒ public ☐ []

Remote IP: 192.168.1.19

Security Model: v3

Security Level: authPriv

User Authentication

Authentication Protocol: MD5

Authentication Password: greenpeace

Data Privacy

Privacy Protocol: DES56

Privacy Password: einstien

Apply Revert

To show remote SNMPv3 users:

1. Click Administration, SNMP.
2. Select Configure User from the Step list.
3. Select Show SNMPv3 Remote User from the Action list.

Figure 282: Showing Remote SNMPv3 Users

Administration > SNMP

Step: 6. Configure User Action: Show SNMPv3 Remote User

SNMPv3 Remote User List Total: 1

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

Delete Revert

Specifying Trap Managers Use the Administration > SNMP (Configure Trap) page to specify the host devices to be sent traps and the types of traps to send. Traps indicating status changes are issued by the switch to the specified trap managers. You must specify trap managers so that key events are reported by this switch to your management station (using network management software). You can specify up to five management stations that will receive authentication failure messages and other trap messages from the switch.

Command Usage

- ◆ Notifications are issued by the switch as trap messages by default. The recipient of a trap message does not send a response to the switch. Traps are therefore not as reliable as inform messages, which include a request for acknowledgement of receipt. Informs can be used to ensure that critical information is received by the host. However, note that informs consume more system resources because they must be kept in memory until a response is received. Informs also add to network traffic. You should consider these effects when deciding whether to issue notifications as traps or informs.

To send an inform to a SNMPv2c host, complete these steps:

1. Enable the SNMP agent ([page 428](#)).
2. Create a view with the required notification messages ([page 433](#)).
3. Configure the group (matching the community string specified on the Configure Trap - Add page) to include the required notify view ([page 436](#)).
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 428](#)).
2. Create a remote SNMPv3 user to use in the message exchange process ([page 441](#)). 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 433](#)).
4. Create a group that includes the required notify view ([page 436](#)).
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 441](#)), 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 443](#)), 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 283: Configuring Trap Managers (SNMPv1)

Administration > SNMP

Step: 7. Configure Trap Action: Add

IP Address: 192.168.0.3

Version: v1

Community String: private

UDP Port (1-65535): 162

Apply Revert

Figure 284: Configuring Trap Managers (SNMPv2c)

Administration > SNMP

Step: 7. Configure Trap Action: Add

IP Address: 192.168.2.9

Version: v2c

Notification Type: Inform

Timeout (0-2147483647): centiseconds

Retry Times (0-255):

Community String: venus

UDP Port (1-65535):

Apply Revert

Figure 285: Configuring Trap Managers (SNMPv3)

Administration > SNMP

Step: 7. Configure Trap Action: Add

IP Address: 192.168.3.9

Version: v3

Notification Type: Inform

Timeout (0-2147483647): centiseconds

Retry Times (0-255):

Remote User Name: carl

UDP Port (1-65535):

Security Level: noAuthNoPriv

Apply Revert

To show configured trap managers:

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

Figure 286: Showing Trap Managers

	IP Address	Version	Community String/User Name	UDP Port	Security Level	Timeout	Retry Times
☐	192.168.0.3	v2c	venus	162	noAuthNoPriv	1500	3
☐	192.168.3.9	v3	carl	162	noAuthNoPriv	1500	3

Creating SNMP Notification Logs Use the Administration > SNMP (Configure Notify Filter - Add) page to create an SNMP notification log.

Command Usage

- ◆ Systems that support SNMP often need a mechanism for recording Notification information as a hedge against lost notifications, whether there are Traps or Informs that may be exceeding retransmission limits. The Notification Log MIB (NLM, RFC 3014) provides an infrastructure in which information from other MIBs may be logged.
- ◆ Given the service provided by the NLM, individual MIBs can now bear less responsibility to record transient information associated with an event against the possibility that the Notification message is lost, and applications can poll the log to verify that they have not missed any important Notifications.
- ◆ If notification logging is not configured, when the switch reboots, some SNMP traps (such as warm start) cannot be logged.
- ◆ To avoid this problem, notification logging should be configured as described in this section, and these commands stored in the startup configuration file using the System > File (Copy – Running-Config) page as described on [page 79](#). Then when the switch reboots, SNMP traps (such as warm start) can now be logged.
- ◆ Based on the default settings used in RFC 3014, a notification log can contain up to 256 entries, and the entry aging time is 1440 minutes. Information recorded in a notification log, and the entry aging time can only be configured using SNMP from a network management station.
- ◆ When a trap host is created using the Administration > SNMP (Configure Trap – Add) page described on [page 446](#), 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:

First create a trap host using the Administration > SNMP (Configure Trap – Add) page. A default notify filter will be created.

Figure 287: Creating a Trap Host

Administration > SNMP

Step: 7. Configure Trap Action: Show

SNMP Trap Manager List Total: 1

	IP Address	Version	Community String/User Name	UDP Port	Security Level	Timeout	Retry Times
☐	192.168.0.99	v2c	test	162			

Delete Revert

Figure 288: Showing the Notification Filter

Administration > SNMP

Step: 8. Configure Notify Filter Action: Show

SNMP Notify Filter List Total: 1

	Filter profile name	IP Address
☐	1	192.168.0.99

Delete Revert

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 289: Creating SNMP Notification Logs

Administration > SNMP

Step: 8. Configure Notify Filter Action: Add

IP Address: 192.168.0.99


Filter Profile Name: R&D

Apply Revert

To show configured SNMP notification logs:

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

Figure 290: 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 291: Showing SNMP Statistics

SNMP Statistics	
SNMP packets input	0
Bad SNMP version errors	0
Unknown community name	0
Illegal operation for community name supplied	0
Encoding errors	0
Number of requested variables	0
Number of altered variables	0
Get-request PDUs	0
Get-next PDUs	0
Set-request PDUs	0
SNMP packets output	0
Too big errors	0
No such name errors	0
Bad values errors	0
General errors	0
Response PDUs	0
Trap PDUs	0

Remote Monitoring

Remote Monitoring allows a remote device to collect information or respond to specified events on an independent basis. This switch is an RMON-capable device which can independently perform a wide range of tasks, significantly reducing network management traffic. It can continuously run diagnostics and log information on network performance. If an event is triggered, it can automatically notify the network administrator of a failure and provide historical information about the event. If it cannot connect to the management agent, it will continue to perform any specified tasks and pass data back to the management station the next time it is contacted.

The switch supports mini-RMON, which consists of the Statistics, History, Event and Alarm groups. When RMON is enabled, the system gradually builds up information about its physical interfaces, storing this information in the relevant RMON database group. A management agent then periodically communicates with the switch using the SNMP protocol. However, if the switch encounters a critical event, it can automatically send a trap message to the management agent which can then respond to the event if so configured.

Configuring RMON Alarms

Use the Administration > RMON (Configure Global - Add - Alarm) page to define specific criteria that will generate response events. Alarms can be set to test data over any specified time interval, and can monitor absolute or changing values (such as a statistical counter reaching a specific value, or a statistic changing by a certain amount over the set interval). Alarms can be set to respond to rising or falling thresholds. (However, note that after an alarm is triggered it will not be triggered again until the statistical value crosses the opposite bounding threshold and then back across the trigger threshold.

Command Usage

- ◆ If an alarm is already defined for an index, the entry must be deleted before any changes can be made.

Parameters

These parameters are displayed:

- ◆ **Index** – Index to this entry. (Range: 1-65535)
- ◆ **Variable** – The object identifier of the MIB variable to be sampled. Only variables of the type `etherStatsEntry.n.n` may be sampled.

Note that `etherStatsEntry.n` uniquely defines the MIB variable, and `etherStatsEntry.n.n` defines the MIB variable, plus the `etherStatsIndex`. For example, `1.3.6.1.2.1.16.1.1.1.6.1` 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 292: Configuring an RMON Alarm

The screenshot shows the 'Administration > RMON' web interface. At the top, there is a breadcrumb trail. Below it, a 'Step:' dropdown is set to '1. Configure Global' and an 'Action:' dropdown is set to 'Add'. The main form has two radio buttons: 'Alarm' (selected) and 'Event'. Below these are several input fields: 'Index (1-65535)' with value '1', 'Variable' with value '6.1', 'Interval (1-31622400)' with value '15' and a 'sec' unit, 'Sample Type' with a dropdown set to 'Delta', 'Rising Threshold (0-2147483647)' with value '100', 'Rising Event Index (0-65535)' with value '30', 'Falling Threshold (0-2147483647)' with value '1', 'Falling Event Index (0-65535)' with value '2', and 'Owner' with value 'bill'. At the bottom right are 'Apply' and 'Revert' buttons.

Field	Value
Index (1-65535)	1
Variable	6.1
Interval (1-31622400)	15 sec
Sample Type	Delta
Rising Threshold (0-2147483647)	100
Rising Event Index (0-65535)	30
Falling Threshold (0-2147483647)	1
Falling Event Index (0-65535)	2
Owner	bill

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 293: Showing Configured RMON Alarms

Administration > RMON

Step: 1. Configure Global Action: Show

☒ Alarm ☐ Event

RMON Alarm List Total 28

☐	Index	Status	Variable	Interval	Type	Last Value	Rising Threshold	Rising Event Index	Falling Threshold	Falling Event Index	Owner
☐	1	Valid	1.3.6.1.2.1.16.1.1.1.6.1	30	Delta	0	892800	0	446400	0	
☐	2	Valid	1.3.6.1.2.1.16.1.1.1.6.2	30	Delta	0	892800	0	446400	0	
☐	3	Valid	1.3.6.1.2.1.16.1.1.1.6.3	30	Delta	0	892800	0	446400	0	
☐	4	Valid	1.3.6.1.2.1.16.1.1.1.6.4	30	Delta	0	892800	0	446400	0	
☐	5	Valid	1.3.6.1.2.1.16.1.1.1.6.5	30	Delta	0	892800	0	446400	0	

Configuring RMON Events

Use the Administration > RMON (Configure Global - Add - Event) page to set the action to take when an alarm is triggered. The response can include logging the alarm or sending a message to a trap manager. Alarms and corresponding events provide a way of immediately responding to critical network problems.

Command Usage

- ◆ If an alarm is already defined for an index, the entry must be deleted before any changes can be made.
- ◆ One default event is configured as follows:
 - event Index = 1
 - Description: RMON_TRAP_LOG
 - Event type: log & trap
 - Event community name is public
 - Owner is RMON_SNMP

Parameters

These parameters are displayed:

- ◆ **Index** – Index to this entry. (Range: 1-65535)
- ◆ **Type** – Specifies the type of event to initiate:
 - **None** – No event is generated.
 - **Log** – Generates an RMON log entry when the event is triggered. Log messages are processed based on the current configuration settings for event logging (see [“System Log Configuration” on page 396](#)).
 - **Trap** – Sends a trap message to all configured trap managers (see [“Specifying Trap Managers” on page 446](#)).
 - **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 429](#)) 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 294: Configuring an RMON Event

The screenshot shows a web interface titled "Administration > RMON". At the top, there are two dropdown menus: "Step:" with "1. Configure Global" selected, and "Action:" with "Add" selected. Below these are two radio buttons: "Alarm" (unselected) and "Event" (selected). The main configuration area contains several labeled text input fields and a dropdown menu:


- Index (1-65535):** A text box containing the number "2".
- Type:** A dropdown menu with "Log and Trap" selected.
- Community:** A text box containing the word "private".
- Description:** A text box containing the text "for software group".
- Owner:** A text box containing the name "david".

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

To show configured RMON events:

1. Click Administration, RMON.
2. Select Configure Global from the Step list.
3. Select Show from the Action list.
4. Click Event.

Figure 295: Showing Configured RMON Events



The screenshot shows the 'Administration > RMON' page. At the top, there are dropdowns for 'Step: 1. Configure Global' and 'Action: Show'. Below these are radio buttons for 'Alarm' and 'Event', with 'Event' selected. The main section is titled 'RMON Event List Total: 4'. It contains a table with 8 columns: Index, Status, Type, Community, Description, Owner, and Last Fired. There are four rows of event configurations. At the bottom of the table are 'Delete' and 'Revert' buttons.

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

Configuring RMON History Samples

Use the Administration > RMON (Configure Interface - Add - History) page to collect statistics on a physical interface to monitor network utilization, packet types, and errors. A historical record of activity can be used to track down intermittent problems. The record can be used to establish normal baseline activity, which may reveal problems associated with high traffic levels, broadcast storms, or other unusual events. It can also be used to predict network growth and plan for expansion before your network becomes too overloaded.

Command Usage

- ◆ Each index number equates to a port on the switch.
- ◆ If history collection is already enabled on an interface, the entry must be deleted before any changes can be made.
- ◆ The information collected for each sample includes:
input octets, packets, broadcast packets, multicast packets, undersize packets, oversize packets, fragments, jabbers, CRC alignment errors, collisions, drop events, and network utilization.
For a description of the statistics displayed on the Show Details page, refer to ["Showing Port or Trunk Statistics" on page 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. (Range: 1-28/52)
- ◆ **Index** – Index to this entry. (Range: 1-65535)
- ◆ **Interval** – The polling interval. (Range: 1-3600 seconds; Default: 1800 seconds)
- ◆ **Buckets** – The number of buckets requested for this entry. (Range: 1-65536; Default: 8)
The number of buckets granted are displayed on the Show page.
- ◆ **Owner** – Name of the person who created this entry. (Range: 1-32 characters)

Web Interface

To periodically sample statistics on a port:

1. Click Administration, RMON.
2. Select Configure Interface from the Step list.
3. Select Add from the Action list.
4. Click History.
5. Select a port from the list as the data source.
6. Enter an index number, the sampling interval, the number of buckets to use, and the name of the owner for this entry.
7. Click Apply

Figure 296: Configuring an RMON History Sample

The screenshot shows the 'Administration > RMON' web interface. At the top, there is a breadcrumb trail 'Administration > RMON'. Below it, there are two dropdown menus: 'Step: 2. Configure Interface' and 'Action: Add'. Underneath these, there are two radio buttons: 'History' (selected) and 'Statistics'. Below the radio buttons, there is a 'Port' dropdown menu with '2' selected. Below the 'Port' dropdown, there are four input fields: 'Index (1-65535)' with the value '100', 'Interval (1-3600)' with the value '60' and a 'sec' label, 'Buckets (1-65536)' 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 297: Showing Configured RMON History Samples

	Index	Status	Interval	Requested Buckets	Granted Buckets	Owner
☐	1	Valid	1800	50	8	
☐	2	Valid	30	50	8	

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 298: 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	756105	3218	91	894	0	0	0	0	0	0	0	0
2	71	00:35:01	21490	76	0	15	0	0	0	0	0	0	0	0
2	72	00:35:31	46521	120	0	15	0	0	0	0	0	0	0	0
2	73	00:36:01	21682	79	0	15	0	0	0	0	0	0	0	0
2	74	00:36:31	21554	77	0	15	0	0	0	0	0	0	0	0

Configuring RMON Statistical Samples Use the Administration > RMON (Configure Interface - Add - Statistics) page to collect statistics on a port, which can subsequently be used to monitor the network for common errors and overall traffic rates.

Command Usage

- ◆ If statistics collection is already enabled on an interface, the entry must be deleted before any changes can be made.
- ◆ The information collected for each entry includes:
input octets, packets, broadcast packets, multicast packets, undersize packets, oversize packets, CRC alignment errors, jabbers, fragments, collisions, drop events, and frames of various sizes.

Parameters

These parameters are displayed:

- ◆ **Port** – The port number on the switch. (Range: 1-28/52)
- ◆ **Index** – Index to this entry. (Range: 1-65535)
- ◆ **Owner** – Name of the person who created this entry. (Range: 1-127 characters)

Web Interface

To enable regular sampling of statistics on a port:

1. Click Administration, RMON.
2. Select Configure Interface from the Step list.
3. Select Add from the Action list.
4. Click Statistics.
5. Select a port from the list as the data source.
6. Enter an index number, and the name of the owner for this entry
7. Click Apply

Figure 299: Configuring an RMON Statistical Sample

The screenshot shows the 'Administration > RMON' configuration page. The 'Step' dropdown is set to '2. Configure Interface' and the 'Action' dropdown is set to 'Add'. The 'History' radio button is unselected, and the 'Statistics' radio button is selected. The 'Port' dropdown is set to '2'. The 'Index (1-65535)' text input field contains the value '100'. The 'Owner' text input field contains the value 'mary'. At the bottom right, there are 'Apply' and 'Revert' buttons.

To show configured RMON statistical samples:

1. Click Administration, RMON.
2. Select Configure Interface from the Step list.
3. Select Show from the Action list.
4. Select a port from the list.
5. Click Statistics.

Figure 300: Showing Configured RMON Statistical Samples


The screenshot shows the 'Administration > RMON' configuration page. The 'Step' dropdown is set to '2. Configure Interface' and the 'Action' dropdown is set to 'Show'. The 'History' radio button is unselected, and the 'Statistics' radio button is selected. The 'Port' dropdown is set to '2'. Below the form, there is a table titled 'RMON Statistics Port List' with a 'Total: 2' indicator. The table has four columns: a checkbox, 'Index', 'Status', and 'Owner'. It contains two rows of data. At the bottom right, there are 'Delete' and 'Revert' buttons.

	Index	Status	Owner
☐	1	Valid	abc
☐	2	Valid	test

To show collected RMON statistical samples:

1. Click Administration, RMON.
2. Select Configure Interface from the Step list.
3. Select Show Details from the Action list.
4. Select a port from the list.
5. Click Statistics.

Figure 301: 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 163](#)).
 2. Add the participating ports to this VLAN (see [“Adding Static Members to VLANs” on page 165](#)), and set them to hybrid mode, tagged members, PVID = 1, and acceptable frame type = all.
- ◆ After the Commander and Members have been configured, any switch in the cluster can be managed from the web agent by choosing the desired Member ID from the Show Member page.

Configuring General Settings for Clusters

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

Command Usage

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

Parameters

These parameters are displayed:

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

Web Interface

To configure a switch cluster:

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

Figure 302: Configuring a Switch Cluster

The screenshot shows the 'Administration > Cluster' configuration page. At the top, there's a breadcrumb 'Administration > Cluster'. Below it, a 'Step:' dropdown menu is set to '1. Configure Global'. The main configuration area contains several fields: 'Cluster Status' with a checked 'Enabled' checkbox, 'Commander Status' with a checked 'Enabled' checkbox, 'IP Pool' with a text input field containing '10.254.254.1', 'Role' with a dropdown menu set to 'Commander', 'Number of Members' with a text input field containing '2', and 'Number of Candidates' with a text input field containing '3'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

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

Parameters

These parameters are displayed:

- ◆ **Member ID** – Specify a Member ID number for the selected Candidate switch. (Range: 1-16)
- ◆ **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 303: Configuring a Cluster Members

Administration > Cluster

Step: 2. Configure Member Action: Add

Member ID (1-36)

MAC Address

☒ Candidate 11-22-33-44-55-11

(xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx)

Apply

Revert

To show the cluster members:

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

Figure 304: Showing Cluster Members

Administration > Cluster

Step: 2. Configure Member Action: Show

Cluster Member List Total: 2

	Member ID	Role	IP Address	MAC Address	Description
☐	1	Active Member	10.254.254.2	11-22-33-44-55-33	ECS4100-26TX
☐	2	Candidate	10.254.254.3	11-22-33-44-55-77	ECS4100-26TX

Delete

Revert

To show cluster candidates:

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

Figure 305: Showing Cluster Candidates

Administration > Cluster

Step: 2. Configure Member Action: Show Candidate

Cluster Candidate List Total: 4

Role	MAC Address	Description
Candidate	11-22-33-44-55-11	ECS4100-26TX
Active Member	11-22-33-44-55-22	ECS4100-26TX
Candidate	11-22-33-44-55-33	ECS4100-26TX
Candidate	11-22-33-44-55-44	ECS4100-26TX

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 306: Managing a Cluster Member



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 307: Setting the Name of a Time Range

Administration > Time Range

Action: Add ▾

Time-Range Name

Apply Revert

To show a list of time ranges:

1. Click Administration, Time Range.
2. Select Show from the Action list.

Figure 308: Showing a List of Time Ranges

Administration > Time Range

Action: Show ▾

Time-Range List Total: 1

	Time-Range Name
☐	r&d

Delete 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 309: Add a Rule to a Time Range

To show the rules configured for a time range:

1. Click Administration, Time Range.
2. Select Show Rule from the Action list.

Figure 310: Showing the Rules Configured for a Time Range

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 than 500 ms.

Operational Concept

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

Ring nodes may be in one of two states:

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

Protection – Protection switching in effect after identifying a signal fault


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

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

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

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

Figure 311: 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 312 on page 474 (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 312 on page 474](#) (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 312: Ring Interconnection Architecture (Multi-ring/Ladder Network)



Configuration Guidelines for ERPS

1. Create an ERPS ring ([Configure Domain – Add](#)): The ring name is used as an index in the G.8032 database.
2. Configure the east and west interfaces ([Configure Domain – Configure Details](#)): Each node on the ring connects to it through two ring ports. Configure one port connected to the next node in the ring to the east (or clockwise direction) and another port facing west in the ring.
3. Configure the RPL owner ([Configure Domain – Configure Details](#)): Configure one node in the ring as the Ring Protection Link (RPL) owner. When this switch is configured as the RPL owner, the west ring port is set as being connected to the RPL. Under normal operations (Idle state), the RPL is blocked to ensure that a loop cannot form in the ring. If a signal failure brings down any other link in the ring, the RPL will be unblocked (Protection state) to ensure proper connectivity among all ring nodes until the failure is recovered.
4. Configure ERPS timers ([Configure Domain – Configure Details](#)): Set the Guard timer to prevent ring nodes from receiving outdated R-APS messages, the

Hold-off timer to filter out intermittent link faults, and the WTR timer to verify that the ring has stabilized before blocking the RPL after recovery from a signal failure.

5. Configure the ERPS control VLAN ([Configure Domain – Configure Details](#)): Specify the control VLAN (CVLAN) used to pass R-APS ring maintenance commands. The CVLAN must NOT be configured with an IP address. In addition, only ring ports may be added to the CVLAN (prior to configuring the VLAN as a CVLAN). No other ports can be members of this VLAN (once set as a CVLAN). Also, the ring ports of the CVLAN must be tagged. Failure to observe these restrictions can result in a loop in the network.
6. Enable ERPS ([Configure Global](#)): Before enabling a ring as described in the next step, first globally enable ERPS on the switch. If ERPS has not yet been enabled or has been disabled, no ERPS rings will work.
7. Enable an ERPS ring ([Configure Domain – Configure Details](#)): Before an ERPS ring can work, it must be enabled. When configuration is completed and the ring enabled, R-APS messages will start flowing in the control VLAN, and normal traffic will begin to flow in the data VLANs. A ring can be stopped by disabling the Admin Status on any node.
8. Display ERPS status information ([Configure Domain – Show](#)): Display ERPS status information for all configured rings.

Configuration Limitations for ERPS

The following configuration limitations apply to ERPS:

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

ERPS Global Configuration Use the Administration > ERPS (Configure Global) page to globally enable or disable ERPS on the switch.

Parameters

These parameters are displayed:

- ◆ **ERPS Status** – Enables ERPS on the switch. (Default: Disabled)

ERPS must be enabled globally on the switch before it can be enabled on an ERPS ring (by setting the Admin Status on the [Configure Domain – Configure Details](#) page).

Web Interface

To globally enable ERPS on the switch:

1. Click Administration, ERPS.
2. Select Configure Global from the Step list.
3. Mark the ERPS Status check box.
4. Click Apply.

Figure 313: Setting ERPS Global Status

The screenshot shows the 'Administration > ERPS' web interface. At the top, there is a breadcrumb trail 'Administration > ERPS'. Below it, a 'Step:' dropdown menu is set to '1. Configure Global'. The main configuration area contains a label 'ERPS Status' followed by a checked checkbox and the text 'Enabled'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

ERPS Ring Configuration Use the Administration > ERPS (Configure Domain) pages to configure ERPS rings.

Command Usage

Ring Initialization

An ERPS ring containing one Control VLAN and one or more protected Data VLANs must be configured, and the global ERPS function enabled on the switch (see [“ERPS Global Configuration” on page 476](#)) before a ring can start running. Once enabled, the RPL owner node and non-owner node state machines will start, and the ring will enter the active state.

Limitations

When configuring a ring port, note that these ports cannot be part of a spanning tree, nor can they be members of a static or dynamic trunk.

Parameters

These parameters are displayed:

Add

- ◆ **Domain Name** – Name of an ERPS ring. (Range: 1-12 characters)
- ◆ **Domain ID** – ERPS ring identifier used in R-APS messages. (Range: 1-255)

Show

- ◆ **Domain Name** – Name of a configured ERPS ring.
- ◆ **ID** – ERPS ring identifier used in R-APS messages.
- ◆ **Admin Status** – Shows whether ERPS is enabled on the switch.
- ◆ **Ver** – Shows the ERPS version.
- ◆ **MEG Level** – The maintenance entity group (MEG) level providing a communication channel for ring automatic protection switching (R-APS) information.
- ◆ **Control VLAN** – Shows the Control VLAN ID.
- ◆ **Node State** – Shows the following ERPS states:
 - **Init** – The ERPS ring has started but has not yet determined the status of the ring.
 - **Idle** – If all nodes in a ring are in this state, it means that all the links in the ring are up. This state will switch to protection state if a link failure occurs.
 - **Protection** – If a node in this state, it means that a link failure has occurred. This state will switch to idle state if all the failed links recover.
- ◆ **Type** – Shows node type as None, RPL Owner or RPL Neighbor.
- ◆ **Revertive** – Shows if revertive or non-revertive recovery is selected.
- ◆ **W/E** – Shows information on the west and east ring port for this node.
- ◆ **West Port** – Shows the west ring port for this node.
- ◆ **East Port** – Shows the east ring port for this node.
- ◆ **Interface** – The port or trunk which is configured as a ring port.
- ◆ **Port State** – The operational state:
 - **Blocking** – The transmission and reception of traffic is blocked and the forwarding of R-APS messages is blocked, but the transmission of locally

generated R-APS messages is allowed and the reception of all R-APS messages is allowed.

- Forwarding – The transmission and reception of traffic is allowed; transmission, reception and forwarding of R-APS messages is allowed.
- Unknown – The interface is not in a known state (includes the domain being disabled).

- ◆ **Local SF** – A signal fault generated on a link to the local node.
- ◆ **Local FS** – Shows if a forced switch command was issued on this interface.
- ◆ **Local MS** – Shows if a manual switch command was issued on this interface.
- ◆ **MEP** – The CFM MEP used to monitor the status on this link.
- ◆ **RPL** – Shows if this node is connected to the RPL.

Configure Details

- ◆ **Domain Name** – Name of a configured ERPS ring. (Range: 1-12 characters)

Service Instances within each ring are based on a unique maintenance association for the specific users, distinguished by the ring name, maintenance level, maintenance association's name, and assigned VLAN. Up to 6 ERPS rings can be configured on the switch.
- ◆ **Domain ID** – ERPS ring identifier used in R-APS messages. (Range: 1-255; Default: None)

R-APS information is carried in an R-APS PDUs. The last octet of the MAC address is designated as the Ring ID (01-19-A7-00-00-[Ring ID]). If use of the default MAC address is disabled for the R-APS Def MAC parameter, then the Domain ID will be used in R-APS PDUs.
- ◆ **Admin Status** – Activates the current ERPS ring. (Default: Disabled)

Before enabling a ring, the global ERPS function should be enabled see ("[ERPS Global Configuration](#)" on page 476), the east and west ring ports configured on each node, the RPL owner specified, and the control VLAN configured.

Once enabled, the RPL owner node and non-owner node state machines will start, and the ring will enter idle state if no signal failures are detected.
- ◆ **Version** – Specifies compatibility with the following ERPS versions:
 - 1 - ERPS version 1 based on ITU-T G.8032/Y.1344.
 - 2 - ERPS version 2 based on ITU-T G.8032/Y.1344 Version 2. (This is the default setting.)

In addition to the basic features provided by version 1, version 2 also supports:

- Multi-ring/ladder network support

- Revertive/Non-revertive recovery
- Forced Switch (FS) and Manual Switch (MS) commands for manually blocking a particular ring port
- Flush FDB (forwarding database) logic which reduces amount of flush FDB operations in the ring
- Support of multiple ERP instances on a single ring

Version 2 is backward compatible with Version 1. If version 2 is specified, the inputs and commands are forwarded transparently. If set to version 1, MS and FS operator commands are filtered, and the switch set to revertive mode.

The version number is automatically set to "1" when a ring node, supporting only the functionalities of G.8032v1, exists on the same ring with other nodes that support G.8032v2.

When ring nodes running G.8032v1 and G.8032v2 co-exist on a ring, the ring ID of each node is configured as "1".

In version 1, the MAC address 01-19-A7-00-00-01 is used for the node identifier. The R-APS Def MAC parameter has no effect.

- ◆ **MEG Level** – The maintenance entity group (MEG) level which provides a communication channel for ring automatic protection switching (R-APS) information. (Range: 0-7)

This parameter is used to ensure that received R-APS PDUs are directed for this ring. A unique level should be configured for each local ring if there are many R-APS PDUs passing through this switch.

- ◆ **Control VLAN** – A dedicated VLAN used for sending and receiving E-APS protocol messages. (Range: 1-4094)

Configure one control VLAN for each ERPS ring. First create the VLAN to be used as the control VLAN (see ["Configuring VLAN Groups" on page 163](#)), add the ring ports for the east and west interface as tagged members to this VLAN (see ["Adding Static Members to VLANs" on page 165](#)), and then use this parameter to add it to the ring.

The following restrictions are recommended to avoid creating a loop in the network or other problems which may occur under some situations:

- The Control VLAN must not be configured as a Layer 3 interface (with an IP address), a dynamic VLAN (with GVRP enabled), nor as a private VLAN.
- In addition, only ring ports may be added to the Control VLAN. No other ports can be members of this VLAN.
- Also, the ring ports of the Control VLAN must be tagged.

Once the ring has been activated, the configuration of the control VLAN cannot be modified. Use the Admin Status parameter to stop the ERPS ring before making any configuration changes to the control VLAN.

- ◆ **Node State** – Refer to the parameters for the Show page.
- ◆ **Node Type** – Shows ERPS node type as one of the following:
 - **None** – Node is neither Ring Protection Link (RPL) owner nor neighbor. (This is the default setting.)
 - **RPL Owner** – Specifies a ring node to be the RPL owner.
 - Only one RPL owner can be configured on a ring. The owner blocks traffic on the RPL during Idle state, and unblocks it during Protection state (that is, when a signal fault is detected on the ring or the protection state is enabled with the Forced Switch or Manual Switch commands on the Configure Operation page).
 - The east and west connections to the ring must be specified for all ring nodes. When this switch is configured as the RPL owner, the west ring port is automatically set as being connected to the RPL.
 - **RPL Neighbor** – Specifies a ring node to be the RPL neighbor.
 - The RPL neighbor node, when configured, is a ring node adjacent to the RPL that is responsible for blocking its end of the RPL under normal conditions (i.e., the ring is established and no requests are present in the ring) in addition to the block at the other end by the RPL Owner Node. The RPL neighbor node may participate in blocking or unblocking its end of the RPL, but is not responsible for activating the reversion behavior.
 - Only one RPL owner can be configured on a ring. If the switch is set as the RPL owner for an ERPS domain, the west ring port is set as one end of the RPL. If the switch is set as the RPL neighbor for an ERPS domain, the east ring port is set as the other end of the RPL.
 - The east and west connections to the ring must be specified for all ring nodes. When this switch is configured as the RPL neighbor, the east ring port is set as being connected to the RPL.
 - Note that is not mandatory to declare a RPL neighbor.
- ◆ **Revertive** – Sets the method of recovery to Idle State through revertive or non-revertive mode. (Default: Enabled)
 - Revertive behavior allows the switch to automatically return the RPL from Protection state to Idle state through the exchange of protocol messages. Non-revertive behavior for Protection, Forced Switch (FS), and Manual Switch (MS) states are basically the same. Non-revertive behavior requires the RPL to be restored from Protection state to Idle state using the Clear command (Configure Operation page).

- Recovery for Protection Switching – A ring node that has one or more ring ports in an SF (Signal Fail) condition, upon detecting the SF condition cleared, keeps at least one of its ring ports blocked for the traffic channel and for the R-APS channel, until the RPL is blocked as a result of ring protection reversion, or until there is another higher priority request (e.g., an SF condition) in the ring.

A ring node that has one ring port in an SF condition and detects the SF condition cleared, continuously transmits the R-APS (NR – no request) message with its own Node ID as the priority information over both ring ports, informing that no request is present at this ring node and initiates a guard timer. When another recovered ring node (or nodes) holding the link block receives this message, it compares the Node ID information with its own Node ID. If the received R-APS (NR) message has the higher priority, this ring node unblocks its ring ports. Otherwise, the block remains unchanged. As a result, there is only one link with one end blocked.

The ring nodes stop transmitting R-APS (NR) messages when they accept an R-APS (NR, RB – RPL Blocked), or when another higher priority request is received.

- Recovery with Revertive Mode – When all ring links and ring nodes have recovered and no external requests are active, reversion is handled in the following way:
 - a. The reception of an R-APS (NR) message causes the RPL Owner Node to start the WTR (Wait-to-Restore) timer.
 - b. The WTR timer is cancelled if during the WTR period a higher priority request than NR is accepted by the RPL Owner Node or is declared locally at the RPL Owner Node.
 - c. When the WTR timer expires, without the presence of any other higher priority request, the RPL Owner Node initiates reversion by blocking its traffic channel over the RPL, transmitting an R-APS (NR, RB) message over both ring ports, informing the ring that the RPL is blocked, and performing a flush FDB action.
 - d. The acceptance of the R-APS (NR, RB) message causes all ring nodes to unblock any blocked non-RPL link that does not have an SF condition. If it is an R-APS (NR, RB) message without a DNF (do not flush) indication, all ring nodes flush the FDB.
- Recovery with Non-revertive Mode – In non-revertive operation, the ring does not automatically revert when all ring links and ring nodes have recovered and no external requests are active. Non-revertive operation is handled in the following way:
 - a. The RPL Owner Node does not generate a response on reception of an R-APS (NR) messages.
 - b. When other healthy ring nodes receive the NR (Node ID) message, no action is taken in response to the message.

- c.** When the operator issues the Clear command (Configure Operation page) for non-revertive mode at the RPL Owner Node, the non-revertive operation is cleared, the RPL Owner Node blocks its RPL port, and transmits an R-APS (NR, RB) message in both directions, repeatedly.
 - d.** Upon receiving an R-APS (NR, RB) message, any blocking node should unblock its non-failed ring port. If it is an R-APS (NR, RB) message without a DNF indication, all ring nodes flush the FDB.
- Recovery for Forced Switching – A Forced Switch command is removed by issuing the Clear command (Configure Operation page) to the same ring node where Forced Switch mode is in effect. The clear command removes any existing local operator commands, and triggers reversion if the ring is in revertive behavior mode.

The ring node where the Forced Switch was cleared continuously transmits the R-APS (NR) message on both ring ports, informing other nodes that no request is present at this ring node. The ring nodes stop transmitting R-APS (NR) messages when they accept an RAPS (NR, RB) message, or when another higher priority request is received.

If the ring node where the Forced Switch was cleared receives an R-APS (NR) message with a Node ID higher than its own Node ID, it unblocks any ring port which does not have an SF condition and stops transmitting R-APS (NR) message over both ring ports.

- Recovery with revertive mode is handled as follows:
 - a.** The reception of an R-APS (NR) message causes the RPL Owner Node to start the WTB timer.
 - b.** The WTB timer is cancelled if during the WTB period a higher priority request than NR is accepted by the RPL Owner Node or is declared locally at the RPL Owner Node.
 - c.** When the WTB timer expires, in the absence of any other higher priority request, the RPL Owner Node initiates reversion by blocking the traffic channel over the RPL, transmitting an R-APS (NR, RB) message over both ring ports, informing the ring that the RPL is blocked, and flushes the FDB.
 - d.** The acceptance of the R-APS (NR, RB) message causes all ring nodes to unblock any blocked non-RPL that does not have an SF condition. If it is an R-APS (NR, RB) message without a DNF indication, all ring nodes flush their FDB. This action unblocks the ring port which was blocked as a result of an operator command.

- Recovery with non-revertive mode is handled as follows:
 - a. The RPL Owner Node, upon reception of an R-APS(NR) message and in the absence of any other higher priority request does not perform any action.
 - b. Then, after the operator issues the Clear command (Configure Operation page) at the RPL Owner Node, this ring node blocks the ring port attached to the RPL, transmits an R-APS (NR, RB) message on both ring ports, informing the ring that the RPL is blocked, and flushes its FDB.
 - c. The acceptance of the R-APS (NR, RB) message triggers all ring nodes to unblock any blocked non-RPL which does not have an SF condition. If it is an R-APS (NR, RB) message without a DNF indication, all ring nodes flush their FDB. This action unblocks the ring port which was blocked as result of an operator command.
- Recovery for Manual Switching – A Manual Switch command is removed by issuing the Clear command (Configure Operation page) at the same ring node where the Manual Switch is in effect. The clear command removes any existing local operator commands, and triggers reversion if the ring is in revertive behavior mode.

The ring node where the Manual Switch was cleared keeps the ring port blocked for the traffic channel and for the R-APS channel, due to the previous Manual Switch command. This ring port is kept blocked until the RPL is blocked as a result of ring protection reversion, or until there is another higher priority request (e.g., an SF condition) in the ring.

The Ethernet Ring Node where the Manual Switch was cleared continuously transmits the R-APS (NR) message on both ring ports, informing that no request is present at this ring node. The ring nodes stop transmitting R-APS (NR) messages when they accept an RAPS (NR, RB) message, or when another higher priority request is received.

If the ring node where the Manual Switch was cleared receives an R-APS (NR) message with a Node ID higher than its own Node ID, it unblocks any ring port which does not have an SF condition and stops transmitting R-APS (NR) message on both ring ports.

- Recovery with revertive mode is handled as follows:
 - a. The RPL Owner Node, upon reception of an R-APS (NR) message and in the absence of any other higher priority request, starts the WTB timer and waits for it to expire. While the WTB timer is running, any latent R-APS (MS) message is ignored due to the higher priority of the WTB running signal.


For example, a node that has one ring port in SF condition and detects that the condition has been cleared, will continuously transmit R-APS (NR) messages with its own Node ID as priority information over both ring ports, informing its neighbors that no request is present at this node. When another recovered node holding the link blocked receives this message, it compares the Node ID information with its own. If the received R-APS (NR) message has a higher priority, this unblocks its ring ports. Otherwise, the block remains unchanged.

The node identifier may also be used for debugging, such as to distinguish messages when a node is connected to more than one ring.

- ◆ **R-APS with VC** – Configures an R-APS virtual channel to connect two interconnection points on a sub-ring, allowing ERPS protocol traffic to be tunneled across an arbitrary Ethernet network. (Default: Enabled)
 - A sub-ring may be attached to a primary ring with or without a virtual channel. A virtual channel is used to connect two interconnection points on the sub-ring, tunneling R-APS control messages across an arbitrary Ethernet network topology. If a virtual channel is not used to cross the intermediate Ethernet network, data in the traffic channel will still flow across the network, but the all R-APS messages will be terminated at the interconnection points.
 - Sub-ring with R-APS Virtual Channel – When using a virtual channel to tunnel R-APS messages between interconnection points on a sub-ring, the R-APS virtual channel may or may not follow the same path as the traffic channel over the network. R-APS messages that are forwarded over the sub-ring's virtual channel are broadcast or multicast over the interconnected network. For this reason the broadcast/multicast domain of the virtual channel should be limited to the necessary links and nodes. For example, the virtual channel could span only the interconnecting rings or sub-rings that are necessary for forwarding R-APS messages of this sub-ring. Care must also be taken to ensure that the local RAPS messages of the sub-ring being transported over the virtual channel into the interconnected network can be uniquely distinguished from those of other interconnected ring R-APS messages. This can be achieved by, for example, by using separate VIDs for the virtual channels of different sub-rings.

Note that the R-APS virtual channel requires a certain amount of bandwidth to forward R-APS messages on the interconnected Ethernet network where a sub-ring is attached. Also note that the protection switching time of the sub-ring may be affected if R-APS messages traverse a long distance over an R-APS virtual channel.


Figure 314: 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 315: Sub-ring without Virtual Channel



- ◆ **R-APS Def MAC** – Sets the switch's MAC address to be used as the node identifier in R-APS messages. (Default: Enabled)

When ring nodes running ERPSv1 and ERPSv2 co-exist on the same ring, the Ring ID of each ring node must be configured as "1".

If this command is disabled, the following strings are used as the node identifier:

- ERPSv1: 01-19-A7-00-00-01
- ERPSv2: 01-19-A7-00-00-[Ring ID]

- ◆ **Propagate TC** – Enables propagation of topology change messages from a secondary ring to the primary ring. (Default: Disabled)


When a secondary ring detects a topology change, it can pass a message about this event to the major ring. When the major ring receives this kind of message from a secondary ring, it can clear the MAC addresses on its ring ports to help the secondary ring restore its connections more quickly through protection switching.

When the MAC addresses are cleared, data traffic may flood onto the major ring. The data traffic will become stable after the MAC addresses are learned again. The major ring will not be broken, but the bandwidth of data traffic on the major ring may suffer for a short period of time due to this flooding behavior.

- ◆ **Non-ERPS Device Protection** – Sends non-standard health-check packets when an owner node enters protection state without any link down event having been detected through Signal Fault messages. (Default: Disabled)
 - The RPL owner node detects a failed link when it receives R-APS (SF - signal fault) messages from nodes adjacent to the failed link. The owner then enters protection state by unblocking the RPL. However, using this standard recovery procedure may cause a non-ERPS device to become isolated when the ERPS device adjacent to it detects a continuity check message (CCM) loss event and blocks the link between the non-ERPS device and ERPS device.

CCMs are propagated by the Connectivity Fault Management (CFM) protocol as described under [“Connectivity Fault Management” on page 505](#). If the standard recovery procedure were used as shown in the following figure, and node E detected CCM loss, it would send an R-APS (SF) message to the RPL owner and block the link to node D, isolating that non-ERPS device.

Figure 316: Non-ERPS Device Protection



When non-ERPS device protection is enabled on the ring, the ring ports on the RPL owner node and non-owner nodes will not be blocked when signal loss is detected by CCM loss events.

- When non-ERPS device protection is enabled on an RPL owner node, it will send non-standard health-check packets to poll the ring health when it enters the protection state. It does not use the normal procedure of waiting to receive an R-APS (NR - no request) message from nodes adjacent to the recovered link. Instead, it waits to see if the non-standard health-check packets loop back. If they do, indicating that the fault has been resolved, the RPL will be blocked.

After blocking the RPL, the owner node will still transmit an R-APS (NR, RB - ring blocked) message. ERPS-compliant nodes receiving this message flush their forwarding database and unblock previously blocked ports. The ring is now returned to Idle state.

- ◆ **Holdoff Timer** – The hold-off timer is used to filter out intermittent link faults. Faults will only be reported to the ring protection mechanism if this timer expires. (Range: 0-10000 milliseconds, in steps of 100 milliseconds)

In order to coordinate timing of protection switches at multiple layers, a hold-off timer may be required. Its purpose is to allow, for example, a server layer protection switch to have a chance to fix the problem before switching at a client layer.

When a new defect or more severe defect occurs (new Signal Failure), this event will not be reported immediately to the protection switching mechanism if the provisioned hold-off timer value is non-zero. Instead, the hold-off timer will be started. When the timer expires, whether a defect still exists or not, the timer will be checked. If one does exist, that defect will be reported to the protection switching mechanism. The reported defect need not be the same one that started the timer.

- ◆ **Guard Timer** – The guard timer is used to prevent ring nodes from receiving outdated R-APS messages. During the duration of the guard timer, all received R-APS messages are ignored by the ring protection control process, giving time for old messages still circulating on the ring to expire. (Range: 10-2000 milliseconds, in steps of 10 milliseconds)

The guard timer duration should be greater than the maximum expected forwarding delay for an R-APS message to pass around the ring. A side-effect of the guard timer is that during its duration, a node will be unaware of new or existing ring requests transmitted from other nodes.

- ◆ **WTB Timer** – The Wait to Block (WTB) timer is used when clearing Forced Switch (FS) and Manual Switch (MS) commands. As multiple FS commands are allowed to co-exist in a ring, the WTB timer ensures that clearing of a single FS command does not trigger re-blocking of the RPL. When clearing an MS command, the WTB timer prevents the formation of a closed loop due to possible a timing anomaly where the RPL owner node receives an outdated remote MS request during the recovery process.

When recovering from an FS or MS command, the delay timer must be long enough to receive any latent remote FS or MS commands. This delay timer called the WTB timer is defined to be 5 seconds longer than the guard timer.

This is enough time to allow a reporting ring node to transmit two R-APS messages and allow the ring to identify the latent condition.

This delay timer is activated on the RPL owner node. When the relevant delay timer expires, the RPL owner node initiates the reversion process by transmitting an R-APS (NR, RB) message. The delay timer, (i.e., WTR or WTB) is deactivated when any higher priority request preempts this delay timer.

The delay timers (i.e. WTR and WTB) may be started and stopped by the system. A request to start running the delay timer does not restart the delay timer. A request to stop the delay timer stops the delay timer and resets its value. The Clear command (Configure Operation page) can be used to stop the delay timer.

- ◆ **WTR Timer** – The wait-to-restore timer is used to verify that the ring has stabilized before blocking the RPL after recovery from a signal failure. (Range: 5-12 minutes)

If the switch goes into ring protection state due to a signal failure, after the failure condition is cleared, the RPL owner will start the wait-to-restore timer and wait until it expires to verify that the ring has stabilized before blocking the RPL and returning to the Idle (normal operating) state.

- ◆ **WTB Expire** – The time before the wait-to-block timer expires.
- ◆ **WTR Expire** – The time before the wait-to-restore timer expires.
- ◆ **West/East** – Connects to next ring node to the west/east.

Each node must be connected to two neighbors on the ring. For convenience, the ports connected are referred to as east and west ports. Alternatively, the closest neighbor to the east should be the next node in the ring in a clockwise direction, and the closest neighbor to the west should be the next node in the ring in a counter-clockwise direction.

- ◆ **Interface** – The port or trunk attached to the west or east ring port.

Note that a ring port cannot be configured as a member of a spanning tree, a dynamic trunk, or a static trunk.

- ◆ **Port State** – Once configured, this field shows the operational state of the ring ports for this node:
 - **Blocking** – The transmission and reception of traffic is blocked and the forwarding of R-APS messages is blocked, but the transmission of locally generated R-APS messages is allowed and the reception of all R-APS messages is allowed.
 - **Forwarding** – The transmission and reception of traffic is allowed; transmission, reception and forwarding of R-APS messages is allowed.
 - **Unknown** – The interface is not in a known state.
- ◆ **Local SF** – Shows if a signal fault exists on a link to the local node.

- ◆ **Local FS** – Shows if a forced switch command was issued on this interface.
- ◆ **Local MS** – Shows if a manual switch command was issued on this interface.
- ◆ **MEP** – Specifies the CCM MEPs used to monitor the link on a ring node.

If a MEP is used to monitor the link status of an ERPS node with CFM continuity check messages, then the MEG Level parameter on this configuration page must match the authorized maintenance level of the CFM domain to which the specified MEP belongs. (See [“Configuring CFM Maintenance Domains” on page 512.](#))

To ensure complete monitoring of a ring node, specify the CFM MEPs used to monitor both the east and west ports of the ring node.

If CFM determines that a MEP node which has been configured to monitor a ring port with this command has gone down, this information is passed to ERPS, which in turn processes it as a ring node failure. For more information on how ERPS recovers from a node failure, refer to the description of the Revertive parameter on this configuration page.

- ◆ **RPL** – If node is connected to the RPL, this shows by which interface.

Web Interface

To create an ERPS ring:

1. Click Administration, ERPS.
2. Select Configure Domain from the Step list.
3. Select Add from the Action list.
4. Enter a name and optional identifier for the ring.
5. Click Apply.

Figure 317: Creating an ERPS Ring

Administration > ERPS

Step: 2. Configure Domain Action: Add

Domain Name rd1

Domain ID (1-255) ☒ 1

Apply Revert

To configure the ERPS parameters for a ring:

1. Click Administration, ERPS.
2. Select Configure Domain from the Step list.
3. Select Configure Details from the Action list.
4. Configure the ERPS parameters for this node. Note that spanning tree protocol cannot be configured on the ring ports, nor can these ports be members of a static or dynamic trunk. And the control VLAN must be unique for each ring. Adjust the protocol timers as required. The RPL owner must be set on one of the rings. And the administrative status enabled once all of the other settings have been entered.
5. Click Apply.

Figure 318: Creating an ERPS Ring

The screenshot displays the 'Administration > ERPS' configuration page. The 'Step' is set to '2. Configure Domain' and the 'Action' is 'Configure Details'. The configuration parameters are as follows:

Parameter	Value
Domain Name	rd1
Domain ID	1
Admin Status	☒ Enabled
Version	2
MEG Level (0-7)	1
Control VLAN	☒ 2
Node State	Idle
Node Type	RPL Owner
Revertive	☒ Enabled
Major Domain	☐ ☐
Node ID	00-E0-0C-00-00-FD
R-APS with VC	☒ Enabled
R-APS Def MAC	☒ Enabled
Propagate TC	☐ Enabled
Non-ERPS Dev Protect	☐ Enabled
Holdoff Timer (0-10000)	0 ms
Guard Timer (10-2000)	500 ms
WTB Timer	5500 ms
WTR Timer (5-12)	5 min
WTB Expire	
WTR Expire	
West	☒ Enabled
Interface	Eth 1/10
Port State	Blocking
Local SF	No
Local FS	No
Local MS	No
MEP (1-8191)	☐
RPL	Yes
East	☒ Enabled
Interface	Eth 1/12
Port State	Forwarding
Local SF	No
Local FS	No
Local MS	No
MEP (1-8191)	☐
RPL	No

Buttons: Apply, Revert

To show the configured ERPS rings:

1. Click Administration, ERPS.
2. Select Configure Domain from the Step list.
3. Select Show from the Action list.

Figure 319: Showing Configured ERPS Rings

Administration > ERPS

Step: 2. Configure Domain Action: Show

Domain List Total: 1

	Domain Name	ID	Admin Status	Ver	MEG Level	Control VLAN	Node State	Type	Revertive	W/E	Interface	Port State	Local SF	Local FS	Local MS	MEP	RPL
☐	rd1	1	Enabled	2	1	2	Idle	RPL Owner	Yes	West	Eth 1/10	Blocking	No	No	No		Yes
										East	Eth 1/12	Forwarding	No	No	No		No

Delete Revert

ERPS Forced and Manual Mode Operations

Use the Administration > ERPS (Configure Operation) page to block a ring port using Forced Switch or Manual Switch commands.

Parameters

These parameters are displayed:

- ◆ **Domain Name** – Name of a configured ERPS ring.
- ◆ **Operation** – Specifies a Forced Switch (FS) or Manual Switch (MS) operation on the east or west ring port.
 - **Forced Switch** – Blocks specified ring port. (Options: West or East)
 - A ring with no pending request has a logical topology with the traffic channel blocked at the RPL and unblocked on all other ring links. In this situation, the FS command triggers protection switching as follows:
 - a. The ring node where an FS command was issued blocks the traffic channel and R-APS channel on the ring port to which the command was issued, and unblocks the other ring port.
 - b. The ring node where the FS command was issued transmits R-APS messages indicating FS over both ring ports. R-APS (FS) messages are continuously transmitted by this ring node while the local FS command is the ring node's highest priority command (see [Table 35 on page 493](#)). 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 35: 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	

Table 35: ERPS Request/State Priority (Continued)

Request / State and Status	Type	Priority
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 35 on page 493](#)). The R-APS (MS) message informs other ring nodes of the MS command and that the traffic channel is blocked on one ring port.
 - c. If no other higher priority commands exist and assuming the ring node was in Idle state before the manual switch command was issued, the ring node flushes its local FDB.
 - d. A ring node accepting an R-APS (MS) message, without any local higher priority requests unblocks any blocked ring port which does

not have an SF condition. This action subsequently unblocks the traffic channel over the RPL.

- e. A ring node accepting an R-APS (MS) message, without any local higher priority requests stops transmitting R-APS messages.
 - f. A ring node receiving an R-APS (MS) message flushes its FDB.
- Protection switching on a manual switch request is completed when the above actions are performed by each ring node. At this point, traffic flows around the ring are resumed. From this point on, the following rules apply regarding processing of further manual switch commands:
 - a. While an existing manual switch request is present in the ring, any new manual switch request is rejected. The request is rejected at the ring node where the new request is issued and a notification is generated to inform the operator that the new MS request was not accepted.
 - b. A ring node with a local manual switch command which receives an R-APS (MS) message with a different Node ID clears its manual switch request and starts transmitting R-APS (NR) messages. The ring node keeps the ring port blocked due to the previous manual switch command.
 - c. An ring node with a local manual switch command that receives an R-APS message or a local request of higher priority than R-APS (MS) clear its manual switch request. The ring node then processes the new higher priority request.
- Recovery for manual switching under revertive and non-revertive mode is described under the Revertive parameter.
- **Clear** – Manually clears the protection state which has been invoked by a forced switch or manual switch command, and the node is operating under non-revertive mode; or before the WTR or WTB timer expires when the node is operating in revertive mode.
- Two steps are required to make a ring operating in non-revertive mode return to Idle state from forced switch or manual switch state:
 - 1. Issue a Clear command to remove the forced switch command on the node where a local forced switch command is active.
 - 2. Issue a Clear command on the RPL owner node to trigger the reversion.
- The Clear command will also stop the WTR and WTB delay timers and reset their values.
- More detailed information about using this command for non-revertive mode is included under the Revertive parameter. (See

the Command Usage section under “ERPS Ring Configuration” on page 476.)

Web Interface

To block a ring port:

1. Click Administration, ERPS.
2. Select Configure Domain from the Step list.
3. Select Configure Operation from the Action list.
4. Select the domain name from the drop-down list.
5. Specify a Forced Switch, Manual Switch, or Clear operation.
6. Click Apply.

Figure 320: Blocking an ERPS Ring Port

The screenshot shows a web interface for configuring ERPS. At the top, it says "Administration > ERPS". Below this, there are two dropdown menus: "Step:" with "2. Configure Domain" selected, and "Action:" with "Configure Operation" selected. Underneath, there are two more dropdown menus: "Domain Name" with "rd1" selected, and "Operation" with "Clear" selected. At the bottom right, there are two buttons: "Apply" and "Revert".

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: 1-28/52)

- ◆ **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 36: 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 321: Enabling OAM for Local Ports

Administration > OAM > Interface

OAM Port List Total: 28

Port	Admin Status	Operation State	Mode	Critical Link Event		Errored Frame		
				Dying Gasp	Critical Event	Status	Window Size (10-65535 1/10 sec)	Threshold Count (1-65535)
1	☐ Enabled	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
6	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
7	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
8	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
9	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
10	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1

Apply Revert

Displaying Statistics for OAM Messages Use the Administration > OAM > Counters page to display statistics for the various types of OAM messages passed across each port.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28/52)
- ◆ **Clear** – Clears statistical counters for the selected ports.
- ◆ **OAMPDU** – Message types transmitted and received by the OAM protocol, including Information OAMPDU, unique Event OAMPDU, Loopback Control OAMPDU, and Organization Specific OAMPDU.

Web Interface

To display statistics for OAM messages:

1. Click Administration, OAM, Counters.

Figure 322: Displaying Statistics for OAM Messages

Administration > OAM > Counters

OAM Port Counters Total: 28

	Port	OAMPDU							
		Information		Event Notification		Loopback Control		Organization Specific	
		Transmitted	Received	Transmitted	Received	Transmitted	Received	Transmitted	Received
☐	1	3	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 323: 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: 1-28/52)
- ◆ **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 324: Displaying Status of Remote Interfaces



The screenshot shows a web interface titled "Administration > OAM > Remote Interface". Below the title bar, there is a section labeled "OAM Remote Port List" with a "Total: 28" indicator. A table with 7 columns and 11 rows is displayed. The columns are: Port, MAC Address, OUI, Remote Loopback, Unidirectional Function, Link Monitor, and MIB Variable Retrieval. The first 10 rows of the table are populated with data, and the 11th row is empty. The data in the first 10 rows is as follows:

Port	MAC Address	OUI	Remote Loopback	Unidirectional Function	Link Monitor	MIB Variable Retrieval
1	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
2	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
3	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
4	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
5	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
6	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
7	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
8	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
9	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
10	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled

Configuring a Remote Loopback Test

Use the Administration > OAM > Remote Loopback (Remote Loopback Test) page to initiate a loop back test to the peer device attached to the selected port.

Command Usage

- ◆ You can use this command to perform an OAM remote loop back test on the specified port. The port that you specify to run this test must be connected to a peer OAM device capable of entering into OAM remote loop back mode.
- ◆ During a remote loop back test, the remote OAM entity loops back every frame except for OAMPDUs and pause frames.
- ◆ OAM remote loopback can be used for fault localization and link performance testing. Statistics from both the local and remote DTE can be queried and compared at any time during loop back testing.
- ◆ To perform a loopback test, first enable Remote Loop Back Mode, click Test, and then click End. The number of packets transmitted and received will be displayed.

Parameters

These parameters are displayed:

Loopback Mode of Remote Device

- ◆ **Port** – Port identifier. (Range: 1-28/52)
- ◆ **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 37: 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 325: Running a Remote Loop Back Test

Administration > OAM > Remote Loopback

Action: Remote Loopback Test

Port: 2

Loopback Mode: ☒ Enabled

Loopback Status: Remote Loopback

Apply Revert

Packet Number (1-99999999):

Packet Size (64-1518): bytes

Test End

Result

10 packets transmitted. 10 packets received. Loss rate is 0.00 %

Displaying Results of Remote Loopback Testing

Use the Administration > OAM > Remote Loopback (Show Test Result) page to display the results of remote loop back testing for each port for which this information is available.

Parameters


These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28/52)
- ◆ **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 326: Displaying the Results of Remote Loop Back Testing

Administration > OAM > Remote Loopback

Action:

Port Remote Test Result List Total: 28

Port	Packets Transmitted	Packets Received	Loss Rate
1	0	0	0.00 %
2	0	0	0.00 %
3	0	0	0.00 %
4	0	0	0.00 %
5	0	0	0.00 %
6	0	0	0.00 %
7	0	0	0.00 %
8	0	0	0.00 %
9	0	0	0.00 %
10	0	0	0.00 %

Connectivity Fault Management

Connectivity Fault Management (CFM) is an OAM protocol that includes proactive connectivity monitoring using continuity check messages, fault verification through loop back messages, and fault isolation by examining end-to-end connections between provider edge devices or between customer edge devices.

CFM is implemented as a service level protocol based on service instances which encompass only that portion of the metropolitan area network supporting a specific customer. CFM can also provide controlled management access to a hierarchy of maintenance domains (such as the customer, service provider, and equipment operator).

This switch supports functions for defining the CFM structure, including domains, maintenance associations, and maintenance access points. It also supports fault detection through continuity check messages for all known maintenance points, and cross-check messages which are used to verify a static list of remote maintenance points located on other devices (in the same maintenance association) against those found through continuity check messages. Fault verification is supported using loop back messages, and fault isolation with link trace messages. Fault notification is also provided by SNMP alarms which are automatically generated by maintenance points when connectivity faults or configuration errors are detected in the local maintenance domain.

Key Components of CFM


CFM provides restricted management access to each Service Instance using a structured conceptual network based on these components:

- ◆ A Maintenance Domain defines a part of the network controlled by a single operator, and supports management access to the domain through Domain Service Access Points (DSAPs) configured on the domain boundary, as well as connectivity testing between these DSAPs.

- ◆ A Maintenance Association (MA) contains the DSAPs for an individual Service Instance. DSAPs are the primary maintenance points used to monitor connectivity across a maintenance domain, and are the entry points to the paths which interconnect the access points allocated to a service instance.
- ◆ A Maintenance Level allows maintenance domains to be nested in a hierarchical fashion, providing access to the specific network portions required by each operator. Domains at lower levels may be either hidden or exposed to operators managing domains at a higher level, allowing either course or fine fault resolution.
- ◆ Maintenance End Points (MEPs) which provide full CFM access to a Service Instance (i.e., a specific MA), and Maintenance Intermediate Points (MIPs) which are passive entities that merely validate received CFM messages, or respond to link trace and loop back requests. MIPs are the interconnection points that make up all possible paths between the DSAPs within an MA, and may also include interconnection points in lower-level domains if exposed by CFM settings.


The following figure shows a single Maintenance Domain, with DSAPs located on the domain boundary, and Internal Service Access Points (ISAPs) inside the domain through which frames may pass between the DSAPs.

Figure 327: 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 328: 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 enabled¹⁷ 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 enabled¹⁷, and a CCM is received from all remote MEPs configured in the static list for this maintenance association¹⁸.
- ◆ **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.

17. Cross-checking must be enabled for this type of trap to be reported (see ["Configuring CFM Maintenance Associations"](#)).

18. See ["Configuring Maintenance End Points"](#).

A MEP Missing trap is sent if cross-checking is enabled¹⁷, and no CCM is received for a remote MEP configured in the static list¹⁸.

◆ **Cross Check MEP Unknown** – Sends a trap if an unconfigured MEP comes up.

A MEP Unknown trap is sent if cross-checking is enabled¹⁷, and a CCM is received from a remote MEP that is not configured in the static list¹⁸.

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 329: Configuring Global Settings for CFM

The screenshot displays the 'Administration > CFM' web interface. At the top, there is a 'Step:' dropdown menu set to '1. Configure Global'. Below this, the 'Global Configuration' section contains several settings:

- CFM Status:** A checkbox labeled 'Enabled' which is checked.
- MEP Cross Check Start Delay (1-65535):** A text input field containing '10' followed by 'sec'.
- Link Trace Cache:** A checkbox labeled 'Enabled' which is checked.
- Link Trace Cache Hold Time (1-65535):** A text input field containing '100' followed by 'min'.
- Link Trace Cache Size (1-4095):** A text input field containing '100' followed by 'entries'.

Below the Global Configuration section is the 'SNMP Trap Configuration' section, which lists several traps with checkboxes, all of which are currently unchecked:

- Connectivity Check Config
- Connectivity Check Loop
- Connectivity Check MEP Down
- Connectivity Check MEP Up
- Cross Check MA Up
- Cross Check MEP Missing
- Cross Check MEP Unknown

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 330: Configuring Interfaces for CFM



Administration > CFM	
Steps: 2. Configure Interface	
Interface ☒ Port ☐ Trunk	
Port List Total: 28	
Port	CFM Status
1	☒ Enabled
2	☒ Enabled
3	☒ Enabled
4	☒ Enabled
5	☒ Enabled

Configuring CFM Maintenance Domains Use the Administration > CFM (Configure MD) pages to create and configure a Maintenance Domain (MD) which defines a portion of the network for which connectivity faults can be managed. Domain access points are set up on the boundary of a domain to provide end-to-end connectivity fault detection, analysis, and recovery. Domains can be configured in a hierarchy to provide management access to the same basic network resources for different user levels.

Command Usage

Configuring General Settings

- ◆ Where domains are nested, an upper-level hierarchical domain must have a higher maintenance level than the ones it encompasses. The higher to lower level domain types commonly include entities such as customer, service provider, and operator.
- ◆ More than one domain can be configured at the same maintenance level, but a single domain can only be configured with one maintenance level.
- ◆ If MEPs (see ["Configuring Maintenance End Points"](#)) or MAs (see ["Configuring CFM Maintenance Associations"](#)) are configured for a domain, they must first be removed before you can remove the domain.

Maintenance domains are designed to provide a transparent method of verifying and resolving connectivity problems for end-to-end connections. By default, these connections run between the domain service access points (DSAPs) within each MA defined for a domain, and are manually configured (see ["Configuring Maintenance End Points"](#)).

In contrast, MIPs are interconnection points that make up all possible paths between the DSAPs within an MA. MIPs are automatically generated by the CFM protocol when the MIP Creation Type is set to "Default" or "Explicit," and the MIP creation state machine is invoked (as defined in IEEE 802.1ag). The default option allows MIPs to be created for all interconnection points within an MA, regardless of the domain's level in the maintenance hierarchy (e.g., customer, provider, or operator). While the explicit option only generates MIPs within an MA if its associated domain is not at the bottom of the maintenance hierarchy. This option is used to hide the structure of network at the lowest domain level.

The diagnostic functions provided by CFM can be used to detect connectivity failures between any pair of MEPs in an MA. Using MIPs allows these failures to be isolated to smaller segments of the network.

Allowing the CFM to generate MIPs exposes more of the network structure to users at higher domain levels, but can speed up the process of fault detection and recovery. This trade-off should be carefully considered when designing a CFM maintenance structure.

Also note that while MEPs are active agents which can initiate consistency check messages (CCMs), transmit loop back or link trace messages, and maintain the local CCM database, MIPs, on the other hand, are passive agents which can only validate received CFM messages, and respond to loop back and link trace messages.

The MIP creation method defined for an MA (see ["Configuring CFM Maintenance Associations"](#)) takes precedence over the method defined on the CFM Domain List.

Configuring Fault Notification

- ◆ A fault alarm can generate an SNMP notification. It is issued when the MEP fault notification generator state machine detects that the configured time period (MEP Fault Notify Alarm Time) has passed with one or more defects indicated, and fault alarms are enabled at or above the specified priority level (MEP Fault Notify Lowest Priority). The state machine transmits no further fault alarms until it is reset by the passage of a configured time period (MEP Fault Notify Reset Time) without a defect indication. The normal procedure upon receiving a fault alarm is to inspect the reporting MEP's managed objects using an appropriate SNMP software tool, diagnose the fault, correct it, re-examine the MEP's managed objects to see whether the MEP fault notification generator state machine has been reset, and repeat those steps until the fault is resolved.
- ◆ Only the highest priority defect currently detected is reported in the fault alarm.

Priority levels include the following options:

Table 38: 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 39: 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 331: Configuring Maintenance Domains

The screenshot shows the 'Administration > CFM' configuration page. At the top, there is a 'Step' dropdown menu set to '1. Configure MD' and an 'Action' dropdown menu set to 'Add'. Below these, there are four input fields: 'MD Index (1-65535)' with the value '1', 'MD Name' with the value 'voip', 'MD Level (0-7)' with the value '3', and 'MIP Creation Type' with a dropdown menu set to 'Explicit'. At the bottom right, there are two buttons: 'Apply' and '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 332: Showing Maintenance Domains

The screenshot shows the 'Administration > CFM' configuration page. At the top, there is a 'Step' dropdown menu set to '1. Configure MD' and an 'Action' dropdown menu set to 'Show'. Below these, there is a table titled 'CFM MD List' with a 'Total: 1' indicator. The table has five columns: 'MD Index', 'MD Name', 'MD Level', and 'MIP Creation Type'. There is a checkbox in the first column. The table contains one row with the following values: '1', 'voip', '3', and 'Explicit'. At the bottom right, there are two buttons: 'Delete' and 'Revert'.

	MD Index	MD Name	MD Level	MIP Creation Type
☐	1	voip	3	Explicit

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 333: Configuring Detailed Settings for Maintenance Domains

The screenshot shows a web-based configuration interface for CFM (Connectivity Fault Management). The breadcrumb trail is 'Administration > CFM'. The current step is '3. Configure MD' and the action is 'Configure Details'. The configuration fields are as follows:

Field	Value
MD Index	1
MEP Archive Hold Time (1-65535)	100
MEP Fault Notify Lowest Priority (1-6)	2
MEP Fault Notify Alarm Time (3-10)	3
MEP Fault Notify Reset Time (3-10)	10

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

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 521](#)).
- ◆ An MA must be defined before any associated DSAPs or remote MEPs can be assigned (see [“Configuring Remote Maintenance End Points” on page 523](#)).
- ◆ Multiple domains at the same maintenance level cannot have an MA on the same VLAN (see [“Configuring CFM Maintenance Domains” on page 512](#)).
- ◆ Before removing an MA, first remove the MEPs assigned to it (see [“Configuring Maintenance End Points” on page 521](#)).
- ◆ For a detailed description of the MIP types, refer to the Command Usage section under [“Configuring CFM Maintenance Domains” on page 512](#).

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: 1-43¹⁹ 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)

19. The total length of the MD name and MA name cannot exceed 44 characters.

- ◆ **MA Name Format** – Specifies the name format for the maintenance association as IEEE 802.1ag character based, or ITU-T SG13/SG15 Y.1731 defined ICC-based format.
 - **Character String** – IEEE 802.1ag defined character string format. This is an IETF RFC 2579 DisplayString.
 - **ICC Based** – ITU-T SG13/SG15 Y.1731 defined ICC based format.
- ◆ **Interval Level** – The delay between sending CCMs. The setting for this parameter is expressed as levels 4 through 7, which in turn map to specific intervals of time. (Options: 4 - 1 second, 5 - 10 seconds, 6 - 1 minute, 7 - 10 minutes)
- ◆ **Connectivity Check** – Enables transmission of CCMs. (Default: Disabled)
- ◆ **Cross Check** – Enables cross-checking between a static list of MEPs assigned to other devices within the same maintenance association and the MEPs learned through CCMs.

Before starting the cross-check process, first configure the remote MEPs that exist on other devices inside the maintenance association using the Remote MEP List (see "[Configuring Remote Maintenance End Points](#)"). These remote MEPs are used in the cross-check operation to verify that all endpoints in the specified MA are operational.

The cross-check start delay, which sets the maximum delay this device waits for a remote MEP to come up before starting the cross-check operation, is a domain-level parameter. To set this parameter, use the CFM MD Configuration screen (see "[Configuring CFM Maintenance Domains](#)").
- ◆ **AIS Status** – Enables/disables suppression of the Alarm Indication Signal (AIS). (Default: Disabled)
- ◆ **AIS Period** – Configures the period at which AIS is sent in an MA. (Range: 1 or 60 seconds; Default: 1 second)
- ◆ **AIS Transmit Level** – Configure the AIS maintenance level in an MA. (Range: 0-7; Default is 0)

AIS Level must follow this rule: AIS Level $\geq$ Domain Level
- ◆ **AIS Suppress Alarm** – Enables/disables suppression of the AIS. (Default: Disabled)

Web Interface

To create a maintenance association:

1. Click Administration, CFM.
2. Select Configure MA from the Step list.
3. Select Add from the Action list.

4. Select an entry from the MD Index list.
5. Specify the MAs assigned to each domain, the VLAN through which CFM messages are passed, and the manner in which MIPs can be created within each MA.
6. Click Apply.

Figure 334: Creating Maintenance Associations

Administration > CFM

Step: 4. Configure MA Action: Add

MD Index 1

MA Index (1-2147483647) 1

MA Name rd

Primary VLAN (1-4094) 1

MIP Creation Type Default

Apply Revert

To show the configured maintenance associations:

1. Click Administration, CFM.
2. Select Configure MA from the Step list.
3. Select Show from the Action list.
4. Select an entry from the MD Index list.

Figure 335: Showing Maintenance Associations

Administration > CFM

Step: 4. Configure MA Action: Show

MD Index 1

CFM MA List Total: 1

	MA Index	MA Name	Primary VLAN	MIP Creation Type
☐	1	rd	1	Default

Delete Revert

To configure detailed settings for maintenance associations:

1. Click Administration, CFM.
2. Select Configure MA from the Step list.
3. Select Configure Details from the Action list.

4. Select an entry from MD Index and MA Index.
5. Specify the CCM interval, enable the transmission of connectivity check and cross check messages, and configure the required AIS parameters.
6. Click Apply

Figure 336: Configuring Detailed Settings for Maintenance Associations

The screenshot shows the 'Administration > CFM' configuration page. At the top, there is a breadcrumb 'Administration > CFM'. Below it, a 'Step:' dropdown is set to '4. Configure MA' and an 'Action:' dropdown is set to 'Configure Details'. The main configuration area contains the following settings:

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, 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 337: Configuring Maintenance End Points

The screenshot shows a web interface titled "Administration > CFM". It features a "Step:" dropdown menu set to "3. Configure MEP" and an "Action:" dropdown menu set to "Add". Below these, there are several configuration fields: "MD Index" and "MA Index" are dropdown menus both set to "1"; "MEP ID (1-8191)" is a text input field containing "1"; "MEP Direction" has a checkbox labeled "Up" which is unchecked; and "Interface" has two radio buttons, "Port" (selected) and "Trunk", with a dropdown menu set to "1" next to "Port". At the bottom right, there are "Apply" and "Revert" buttons.

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 338: Showing Maintenance End Points

Administration > CFM

Step: 3. Configure MEP Action: Show

MD Index: 1
MA Index: 1

CFM Maintenance Association End Point List Total: 1

MEP ID	MEP Direction	Interface
1	Down	Unit 1 / Port 1

Delete Revert

Configuring Remote Maintenance End Points

Use the Administration > CFM (Configure Remote MEP – Add) page to specify remote maintenance end points (MEPs) set on other CFM-enabled devices within a common MA. Remote MEPs can be added to a static list in this manner to verify that each entry has been properly configured and is operational. When cross-checking is enabled, the list of statically configured remote MEPs is compared against the MEPs learned through continuity check messages (CCMs), and any discrepancies reported via SNMP traps.

Command Usage

- ◆ All MEPs that exist on other devices inside a maintenance association should be statically configured to ensure full connectivity through the cross-check process.
- ◆ Remote MEPs can only be configured if local domain service access points (DSAPs) have already been created (see ["Configuring Maintenance End Points"](#)) at the same maintenance level and in the same MA. DSAPs are MEPs that exist on the edge of the domain, and act as primary service access points for end-to-end cross-check, loop-back, and link-trace functions.
- ◆ The MEP cross-check start delay which sets the maximum delay that a device waits for remote MEPs to come up before starting the cross-check operation can be configured on the Configure Global page (see ["Configuring Global Settings for CFM"](#)).
- ◆ SNMP traps for continuity check events discovered by cross-check operations can also be configured on the Configure Global page (see ["Configuring Global Settings for CFM"](#)).

Parameters

These parameters are displayed:

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)

- ◆ **MEP ID** – Identifier for a maintenance end point which exists on another CFM-enabled device within the same MA. (Range: 1-8191)

Web Interface

To configure a remote maintenance end point:

1. Click Administration, CFM.
2. Select Configure Remote MEP from the Step list.
3. Select Add from the Action list.
4. Select an entry from MD Index and MA Index.
5. Specify the remote MEPs which exist on other devices within the same MA.
6. Click Apply.

Figure 339: Configuring Remote Maintenance End Points

The screenshot shows a web interface titled "Administration > CFM". Below the title, 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 340: Showing Remote Maintenance End Points

Administration > CFM

Step: 4. Configure Remote MEP Action: Show

MD Index: 1
MA Index: 1

CFM Remote Maintenance Association End Point List Total: 1

	MEP ID
☐	2

Delete Revert

Transmitting Link Trace Messages

Use the Administration > CFM (Transmit Link Trace) page to transmit link trace messages (LTMs). These messages can isolate connectivity faults by tracing the path through a network to the designated target node (i.e., a remote maintenance end point).

Command Usage

- ◆ LTMs can be targeted to MEPs, not MIPs. Before sending a link trace message, be sure you have configured the target MEP for the specified MA (see ["Configuring Remote Maintenance End Points"](#)).
- ◆ If MAC address of target MEP has not been learned by any local MEP, then the link trace may fail. Use the Show Remote MEP page (see ["Displaying Remote MEPs"](#)) to verify that a MAC address has been learned for the target MEP.
- ◆ LTMs are sent as multicast CFM frames, and forwarded from MIP to MIP, with each MIP generating a link trace reply, up to the point at which the LTM reaches its destination or can no longer be forwarded.
- ◆ LTMs are used to isolate faults. However, this task can be difficult in an Ethernet environment, since each node is connected through multipoint links. Fault isolation is even more challenging since the MAC address of the target node can age out in several minutes. This can cause the traced path to vary over time, or connectivity lost if faults cause the target MEP to be isolated from other MEPs in an MA.
- ◆ When using the command line or web interface, the source MEP used by to send a link trace message is chosen by the CFM protocol. However, when using SNMP, the source MEP can be specified by the user.
- ◆ Parameters controlling the link trace cache, including operational state, entry hold time, and maximum size can be configured on the Configure Global page (see ["Configuring Global Settings for CFM"](#)).

Parameters

These parameters are displayed:

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)
- ◆ **Source MEP ID** – The identifier of a source MEP that will send the link trace message. (Range: 1-8191)
- ◆ **Target**
 - **MEP ID** – The identifier of a remote MEP that is the target of a link trace message. (Range: 1-8191)
 - **MAC Address** – MAC address of a remote MEP that is the target of a link trace message. This address can be entered in either of the following formats: xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx
- ◆ **TTL** – The time to live of the link trace message. (Range: 0-255 hops)

Web Interface

To transmit link trace messages:

1. Click Administration, CFM.
2. Select Transmit Link Trace from the Step list.
3. Select an entry from MD Index and MA Index.
4. Specify the source MEP, the target MEP using either its MEP identifier or MAC address, and set the maximum number of hops allowed in the TTL field.
5. Click Apply.
6. Check the results in the Link Trace cache (see ["Displaying the Link Trace Cache"](#)).

Figure 341: Transmitting Link Trace Messages

The screenshot shows a web interface for 'Administration > CFM'. The 'Step' dropdown is set to '7. Transmit Link Trace'. The form contains the following fields:

- MD Index:** A dropdown menu with the value '1' selected.
- MA Index:** A dropdown menu with the value '1' selected.
- Source MEP ID (1-8191):** A text input field containing the value '1'.
- Target:** A section with two radio buttons. The first radio button, labeled 'MEP ID (1-8191)', is selected, and its corresponding text input field contains the value '2'. The second radio button is labeled 'MAC Address' and is unselected, with an empty text input field next to it.
- TTL (0-255):** A text input field containing the value '5'.

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

Transmitting Loopback Messages

Use the Administration > CFM (Transmit Loopback) page to transmit Loopback Messages (LBMs). These messages can be used to isolate or verify connectivity faults by submitting a request to a target node (i.e., a remote MEP or MIP) to echo the message back to the source.

Command Usage

- ◆ Loopback messages can be used for fault verification and isolation after automatic detection of a fault or receipt of some other error report. Loopback messages can also be used to confirm the successful restoration or initiation of connectivity. The receiving maintenance point should respond to the loop back message with a loopback reply.
- ◆ The point from which the loopback message is transmitted (i.e., a local DSAP) and the target maintenance point must be within the same MA.
- ◆ If the continuity check database does not have an entry for the specified maintenance point, an error message will be displayed.
- ◆ When using the command line or web interface, the source MEP used by to send a loopback message is chosen by the CFM protocol. However, when using SNMP, the source MEP can be specified by the user.

Parameters

These parameters are displayed:

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)
- ◆ **Source MEP ID** – The identifier of a source MEP that will send the loopback message. (Range: 1-8191)

◆ **Target**

- **MEP ID** – The identifier of a remote MEP that is the target of a loopback message. (Range: 1-8191)
- **MAC Address** – MAC address of a remote MEP that is the target of a loopback message. This address can be entered in either of the following formats: xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx

◆ **Count** – The number of times the loopback message is sent. (Range: 1-1024)


◆ **Packet Size** – The size of the loopback message. (Range: 64-1518 bytes; Default: 64 bytes)

Web Interface

To transmit loopback messages:

1. Click Administration, CFM.
2. Select Transmit Loopback from the Step list.
3. Select an entry from MD Index and MA Index.
4. Specify the source MEP, the target MEP using either its MEP identifier or MAC address, set the number of times the loopback message is to be sent.
5. Click Apply.

Figure 342: Transmitting Loopback Messages



Administration > CFM

Step: 8. Transmit Loopback

MD Index: 1

MA Index: 1

Source MEP ID (1-8191):

Target: ☒ MEP ID (1-8191)

☐ MAC Address

Counts (1-1024):

Packet Size (64-1518): bytes

Test

Result

```
Sending 200 Ethernet CFM loopback message, timeout is 5 sec.
Received 0 Ethernet CFM loopback message in 1 sec.
Received 0 Ethernet CFM loopback message in 5 secs. Success rate is 0% (0/5).
```

Transmitting Delay-Measure Requests

Use the Administration > CFM (Transmit Delay Measure) page to send periodic delay-measure requests to a specified MEP within a maintenance association.

Command Usage

- ◆ Delay measurement can be used to measure frame delay and frame delay variation between MEPs.
- ◆ A local MEP must be configured for the same MA before you can use this function.
- ◆ If a MEP is enabled to generate frames with delay measurement (DM) information, it periodically sends DM frames to its peer MEP in the same MA., and expects to receive DM frames back from it.
- ◆ Frame delay measurement can be made only for two-way measurements, where the MEP transmits a frame with DM request information with the TxTimeStampf (Timestamp at the time of sending a frame with DM request information), and the receiving MEP responds with a frame with DM reply information with TxTimeStampf copied from the DM request information, RxTimeStampf (Timestamp at the time of receiving a frame with DM request information), and TxTimestampb (Timestamp at the time of transmitting a frame with DM reply information):

$$\text{Frame Delay} = (\text{RxTimeStampb} - \text{TxTimeStampf}) - (\text{TxTimestampb} - \text{RxTimeStampf})$$
- ◆ The MEP can also make two-way frame delay variation measurements based on its ability to calculate the difference between two subsequent two-way frame delay measurements.

Parameters

These parameters are displayed:

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)
- ◆ **Source MEP ID** – The identifier of a source MEP that will send the delay-measure message. (Range: 1-8191)
- ◆ **Target**
 - **MEP ID** – The identifier of a remote MEP that is the target of a delay-measure message. (Range: 1-8191)
 - **MAC Address** – MAC address of a remote MEP that is the target of a delay-measure message. This address can be entered in either of the following formats: xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx
- ◆ **Count** – The number of times to retry sending the message if no response is received before the specified timeout. (Range: 1-5; Default: 5)

- ◆ **Packet Size** – The size of the delay-measure message. (Range: 64-1518 bytes; Default: 64 bytes)
- ◆ **Interval** – The transmission delay between delay-measure messages. (Range: 1-5 seconds; Default: 1 second)
- ◆ **Timeout** – The timeout to wait for a response. (Range: 1-5 seconds; Default: 5 seconds)

Web Interface

To transmit delay-measure messages:

1. Click Administration, CFM.
2. Select Transmit Delay Measure from the Step list.
3. Select an entry from MD Index and MA Index.
4. Specify the source MEP, the target MEP using either its MEP identifier or MAC address, set the number of times the delay-measure message is to be sent, the interval, and the timeout.
5. Click Apply.

Figure 343: Transmitting Delay-Measure Messages

Administration > CFM

Step: 7. Transmit Delay Measure

MD Index: 1
MA Index: 3

Source MEP ID (1-8191):
Target: ☒ MEP ID (1-8191)
☐ MAC Address

Counts (1-5):
Packet Size (64-1518): bytes
Interval (1-5): ms
Timeout (1-5): ms

Apply Revert

Result

Sending 5 Ethernet CFM delay measurement message, timeout is 5 sec.

Sequence	Delay Time (ms.)	Delay Variation (ms.)
1	< 10	0
2	10	10
3	< 10	10
4	< 10	0
5	< 10	0

Success rate is 100% (5/5), delay time min/avg/max=0/2/10 ms.
Average frame delay variation is 4 ms.

Displaying Local MEPs Use the Administration > CFM > Show Information (Show Local MEP) page to show information for the MEPs configured on this device.

Parameters

These parameters are displayed:

- ◆ **MEP ID** – Maintenance end point identifier.
- ◆ **MD Name** – Maintenance domain name.
- ◆ **Level** – Authorized maintenance level for this domain.
- ◆ **Direction** – Direction in which the MEP communicates CFM messages:
 - Down indicates that the MEP is facing away from the switch, and transmits CFM messages towards, and receives them from, the direction of the physical medium.
 - Up indicates that the MEP faces inward toward the switch cross-connect matrix, and transmits CFM messages towards, and receives them from, the direction of the internal bridge relay mechanism.
- ◆ **Primary VLAN** – Service VLAN ID.
- ◆ **Interface** – Physical interface of this entry (either a port or trunk).
- ◆ **CC Status** – Shows administrative status of CCMs.
- ◆ **MAC Address** – MAC address of this MEP entry.

Web Interface

To show information for the MEPs configured on this device:

1. Click Administration, CFM.
2. Select Show Information from the Step list.
3. Select Show Local MEP from the Action list.

Figure 344: Showing Information on Local MEPs

Administration > CFM

Step: 10: Show Information Action: Show Local MEP

CFM Local Maintenance Association End Point Information Total: 1

MEP ID	MD Name	Level	Direction	Primary VLAN	Interface	CC Status	MAC Address
1	voip	3	Down	1	Unit 1 / Port 1	Enabled	00-E0-0C-00-00-FE

Displaying Details for Local MEPs Use the Administration > CFM > Show Information (Show Local MEP Details) page to show detailed CFM information about a local MEP in the continuity check database.

Parameters

These parameters are displayed:

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)
- ◆ **MEP ID** – Maintenance end point identifier. (Range: 1-8191)
- ◆ **MD Name** – The maintenance domain for this entry.
- ◆ **MA Name** – Maintenance association to which this remote MEP belongs.
- ◆ **MA Name Format** – The format of the Maintenance Association name, including Character String or ICC Based.
- ◆ **Level** – Maintenance level of the local maintenance point.
- ◆ **Direction** – The direction in which the MEP faces on the Bridge port (up or down).
- ◆ **Interface** – The port to which this MEP is attached.
- ◆ **CC Status** – Shows if the MEP will generate CCM messages.
- ◆ **MAC Address** – MAC address of the local maintenance point. (If a CCM for the specified remote MEP has never been received or the local MEP record times out, the address will be set to the initial value of all Fs.)
- ◆ **Defect Condition** – Shows the defect detected on the MEP.
- ◆ **Received RDI** – Receive status of remote defect indication (RDI) messages on the MEP.
- ◆ **AIS Status** – Shows if MEPs within the specified MA are enabled to send frames with AIS information following detection of defect conditions.
- ◆ **AIS Period** – The interval at which AIS information is sent.
- ◆ **AIS Transmit Level** – The maintenance level at which AIS information will be sent for the specified MEP.
- ◆ **Suppress Alarm** – Shows if the specified MEP is configured to suppress sending frames containing AIS information following the detection of defect conditions.

- ◆ **Suppressing Alarms** – Shows if the specified MEP is currently suppressing sending frames containing AIS information following the detection of defect conditions.

Web Interface

To show detailed information for the MEPs configured on this device:

1. Click Administration, CFM.
2. Select Show Information from the Step list.
3. Select Show Local MEP Details from the Action list.
4. Select an entry from MD Index and MA Index.
5. Select a MEP ID.

Figure 345: Showing Detailed Information on Local MEPs

Administration > CFM

Step: 10. Show Information
Action: Show Local MEP Details

MD Index 1
MA Index 1
MEP ID 1

Query

MD Name	md1
MA Name	ma1
MA Name Format	Character String
Level	0
Direction	Up
Interface	Unit 1 / Port 2
CC Status	Enabled
MAC Address	00-1B-DS-50-91-FD
Defect Condition	No Defect
Received RDI	False
AIS Status	Enabled
AIS Period	60 sec
AIS Transmit Level	Default
Suppress Alarm	Enabled
Suppressing Alarms	Disabled

Displaying Local MIPs Use the Administration > CFM > Show Information (Show Local MIP) page to show the MIPs on this device discovered by the CFM protocol. (For a description of MIPs, refer to the Command Usage section under "[Configuring CFM Maintenance Domains](#)".)

Parameters

These parameters are displayed:

- ◆ **MD Name** – Maintenance domain name.
- ◆ **Level** – Authorized maintenance level for this domain.
- ◆ **MA Name** – Maintenance association name.
- ◆ **Primary VLAN** – Service VLAN ID.
- ◆ **Interface** – Physical interface of this entry (either a port or trunk).

Web Interface

To show information for the MIPs discovered by the CFM protocol:

1. Click Administration, CFM.
2. Select Show Information from the Step list.
3. Select Show Local MIP from the Action list.

Figure 346: Showing Information on Local MIPs

Administration > CFM

Step: 10. Show Information Action: Show Local MIP

CFM Local Maintenance Association Intermediate Point Information Total: 27

MD Name	Level	MA Name	Primary VLAN	Interface
voip	3	rd	1	Unit 1 / Port 2
voip	3	rd	1	Unit 1 / Port 3
voip	3	rd	1	Unit 1 / Port 4
voip	3	rd	1	Unit 1 / Port 5
voip	3	rd	1	Unit 1 / Port 6
voip	3	rd	1	Unit 1 / Port 7
voip	3	rd	1	Unit 1 / Port 8
voip	3	rd	1	Unit 1 / Port 9
voip	3	rd	1	Unit 1 / Port 10

Displaying Remote MEPs Use the Administration > CFM > Show Information (Show Remote MEP) page to show MEPs located on other devices which have been discovered through continuity check messages, or statically configured in the MEP database and verified through cross-check messages.

Parameters

These parameters are displayed:

- ◆ **MEP ID** – Maintenance end point identifier.
- ◆ **MA Name** – Maintenance association name.
- ◆ **Level** – Authorized maintenance level for this domain.
- ◆ **Primary VLAN** – Service VLAN ID.
- ◆ **MEP Up** – Indicates whether or not this MEP is functioning normally.
- ◆ **Remote MAC Address** – MAC address of the remote maintenance point. (If a CCM for the specified remote MEP has never been received or the remote MEP record times out, the address will be set to the initial value of all Fs.)

Web Interface

To show information for remote MEPs:

1. Click Administration, CFM.
2. Select Show Information from the Step list.
3. Select Show Remote MEP from the Action list.

Figure 347: Showing Information on Remote MEPs

Administration > CFM

Step: 10: Show Information Action: Show Remote MEP

CFM Remote Maintenance Association End Point Information Total: 1

MEP ID	MA Name	Level	Primary VLAN	MEP Up	Remote MAC Address
1	ma1	0	2	Yes	00-01-02-03-04-05

Clear

Displaying Details for Remote MEPs Use the Administration > CFM > Show Information (Show Remote MEP Details) page to show detailed information for MEPs located on other devices which have been discovered through continuity check messages, or statically configured in the MEP database and verified through cross-check messages.

Parameters

These parameters are displayed:

- ◆ **MD Index** – Domain index. (Range: 1-65535)
- ◆ **MA Index** – MA identifier. (Range: 1-2147483647)
- ◆ **MEP ID** – Maintenance end point identifier. (Range: 1-8191)
- ◆ **MD Name** – Maintenance domain name.
- ◆ **MA Name** – Maintenance association name.
- ◆ **Level** – Authorized maintenance level for this domain.
- ◆ **MAC Address** – MAC address of this MEP entry.
- ◆ **Primary VLAN** – Service VLAN ID.
- ◆ **Incoming Port** – Port to which this remote MEP is attached.
- ◆ **CC Lifetime** – Length of time to hold messages about this MEP in the CCM database.
- ◆ **Age of Last CC Message** – Length of time the last CCM message about this MEP has been in the CCM database.
- ◆ **Frame Loss** – Percentage of transmitted frames lost.
- ◆ **CC Packet Statistics** – The number of CCM packets received successfully and those with errors.
- ◆ **Port State** – Port states include:
 - Up – The port is functioning normally.
 - Blocked – The port has been blocked by the Spanning Tree Protocol.
 - No port state – Either no CCM has been received, or no port status TLV was received in the last CCM.
- ◆ **Interface State** – Interface states include:
 - No Status – Either no CCM has been received, or no interface status TLV was received in the last CCM.
 - Up – The interface is ready to pass packets.

- Down – The interface cannot pass packets.
- Testing – The interface is in some test mode.
- Unknown – The interface status cannot be determined for some reason.
- Dormant – The interface is not in a state to pass packets but is in a pending state, waiting for some external event.
- Not Present – Some component of the interface is missing.
- isLowerLayerDown – The interface is down due to state of the lower layer interfaces.

◆ **Crosscheck Status** – Shows if crosscheck function has been enabled.

Web Interface

To show detailed information for remote MEPs:

1. Click Administration, CFM.
2. Select Show Information from the Step list.
3. Select Show Remote MEP Details from the Action list.
4. Select an entry from MD Index and MA Index.
5. Select a MEP ID.

Figure 348: Showing Detailed Information on Remote MEPs

Administration > CFM	
Step:	10. Show Information
Action:	Show Remote MEP Details
MD Index	1
MA Index	1
MEP ID	1
Query	
MD Name	md1
MA Name	ma1
Level	0
MAC Address	00-1B-D5-50-91-FD
Primary VLAN	2
Incoming Port	Unit 1 / Port 2
CC Lifetime	75 sec
Age of Last CC Message	60 sec
Frame Loss	0
CC Packet Statistics	139 / 0 (Received / Error)
Port State	Enabled
Interface State	Enabled
Crosscheck Status	Disabled

Displaying the Link Trace Cache Use the Administration > CFM > Show Information (Show Link Trace Cache) page to show information about link trace operations launched from this device.

Parameters

These parameters are displayed:

- ◆ **Hops** – The number hops taken to reach the target MEP.
- ◆ **MA** – Maintenance association name.
- ◆ **IP 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 349: Showing the Link Trace Cache

Hops	MA	IP/Alias Address	Forwarded	Ingress MAC Address	Egress MAC Address	Ingress Action	Egress Action	Reply
1	denver	11-11-11-11-11	Forwarded	11-11-11-11-12	11-11-11-11-13	OK	OK	FDB
2	boston	AA-BB-CC-00-04-00	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 alarm²⁰.
- ◆ **Alarm Time** – The time a defect must exist before a fault alarm is issued²⁰.
- ◆ **Reset Time** – The time after a fault alarm has been issued, and no defect exists, before another fault alarm can be issued²⁰.

20. See “Configuring CFM Maintenance Domains” on page 512.

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 350: Showing Settings for the Fault Notification Generator

Administration > CFM						
Step: 10: Show Information Action: Show Fault Notification Generator						
CFM Fault Notification Generator Information Total: 1						
MEP ID	MD Name	MA Name	Highest Defect	Lowest Alarm	Alarm Time (sec)	Reset Time (sec)
1	aa	bb	NONE	allDef	3	3

Displaying Continuity Check Errors

Use the Administration > CFM > Show Information (Show Continuity Check Error) page to display the CFM continuity check errors logged on this device.

Parameters

These parameters are displayed:

- ◆ **Level** – Maintenance level associated with this entry.
- ◆ **Primary VLAN** – VLAN in which this error occurred.
- ◆ **MEP ID** – Identifier of remote MEP.
- ◆ **Interface** – Port at which the error was recorded.
- ◆ **Remote MAC** – MAC address of remote MEP.
- ◆ **Reason** – Error types include:
 - **LEAK** – MA x is associated with a specific VID list²¹, 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 list²¹, 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.

21. 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 351: Showing Continuity Check Errors

Administration > CFM

Step: 10. Show Information Action: Show Continuity Check Error

CFM Continuity Check Error Information Total: 1

Level	Primary VLAN	MEP ID	Interface	Remote MAC	Reason	MA Name
5	2	40	Unit 1 / Port 10	00-01-02-03-04-05	LEAK	aa

Clear

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: 1-28/52)

- ◆ **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 40: 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 352: Enabling OAM for Local Ports

Administration > OAM > Interface

OAM Port List Total: 28

Port	Admin Status	Operation State	Mode	Critical Link Event		Errored Frame		
				Dying Gasp	Critical Event	Status	Window Size (10-65535 1/10 sec)	Threshold Count (1-65535)
1	☐ Enabled	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
6	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
7	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
8	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
9	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1
10	☐ Enabled	Disabled	Active	☒ Enabled	☒ Enabled	☒ Enabled	10	1

Apply Revert

Displaying Statistics for OAM Messages Use the Administration > OAM > Counters page to display statistics for the various types of OAM messages passed across each port.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28/52)
- ◆ **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 353: Displaying Statistics for OAM Messages

Administration > OAM > Counters

OAM Port Counters Total: 28

	Port	OAMPDU							
		Information		Event Notification		Loopback Control		Organization Specific	
		Transmitted	Received	Transmitted	Received	Transmitted	Received	Transmitted	Received
☐	1	3	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 354: 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: 1-28/52)
- ◆ **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 355: Displaying Status of Remote Interfaces

Port	MAC Address	OUI	Remote Loopback	Unidirectional Function	Link Monitor	MIB Variable Retrieval
1	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
2	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
3	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
4	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
5	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
6	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
7	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
8	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
9	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled
10	00-00-00-00-00-00	00-00-00	Disabled	Disabled	Disabled	Disabled

Configuring a Remote Loopback Test

Use the Administration > OAM > Remote Loopback (Remote Loopback Test) page to initiate a loop back test to the peer device attached to the selected port.

Command Usage

- ◆ You can use this command to perform an OAM remote loop back test on the specified port. The port that you specify to run this test must be connected to a peer OAM device capable of entering into OAM remote loop back mode.
- ◆ During a remote loop back test, the remote OAM entity loops back every frame except for OAMPDUs and pause frames.
- ◆ OAM remote loopback can be used for fault localization and link performance testing. Statistics from both the local and remote DTE can be queried and compared at any time during loop back testing.
- ◆ To perform a loopback test, first enable Remote Loop Back Mode, click Test, and then click End. The number of packets transmitted and received will be displayed.

Parameters

These parameters are displayed:

Loopback Mode of Remote Device

- ◆ **Port** – Port identifier. (Range: 1-28/52)
- ◆ **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 41: 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 356: Running a Remote Loop Back Test

Administration > OAM > Remote Loopback

Action: Remote Loopback Test

Port: 2

Loopback Mode: ☒ Enabled

Loopback Status: Remote Loopback

Apply Revert

Packet Number (1-99999999):

Packet Size (64-1518): bytes

Test End

Result

10 packets transmitted. 10 packets received. Loss rate is 0.00 %

Displaying Results of Remote Loopback Testing

Use the Administration > OAM > Remote Loopback (Show Test Result) page to display the results of remote loop back testing for each port for which this information is available.

Parameters

These parameters are displayed:


- ◆ **Port** – Port identifier. (Range: 1-28/52)
- ◆ **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 357: Displaying the Results of Remote Loop Back Testing



The screenshot shows a web interface for 'Administration > OAM > Remote Loopback'. It includes an 'Action:' dropdown menu set to 'Show Test Result'. Below this is a 'Port Remote Test Result List' with a 'Total: 28' count. The table displays results for 10 ports, all showing 0 packets transmitted and received, and a 0.00% loss rate. The table has columns for Port, Packets Transmitted, Packets Received, and Loss Rate. There are also pagination controls (1, 2, 3) in the top right corner of the table area.

Port	Packets Transmitted	Packets Received	Loss Rate
1	0	0	0.00 %
2	0	0	0.00 %
3	0	0	0.00 %
4	0	0	0.00 %
5	0	0	0.00 %
6	0	0	0.00 %
7	0	0	0.00 %
8	0	0	0.00 %
9	0	0	0.00 %
10	0	0	0.00 %

UDLD Configuration

The switch can be configured to detect general loopback conditions caused by hardware problems or faulty protocol settings. When enabled, a control frame is transmitted on the participating ports, and the switch monitors inbound traffic to see if the frame is looped back.

Usage Guidelines

- ◆ The default settings for the control frame transmit interval and recover time may be adjusted to improve performance for your specific environment. The shutdown mode may also need to be changed once you determine what kind of packets are being looped back.
- ◆ General loopback detection provided by the commands described in this section and loopback detection provided by the spanning tree protocol cannot both be enabled at the same time. If loopback detection is enabled for the spanning tree protocol, general loopback detection cannot be enabled on the same interface.
- ◆ When a loopback event is detected on an interface or when an interface is released from a shutdown state caused by a loopback event, a trap message is sent and the event recorded in the system log.
- ◆ Loopback detection must be enabled both globally and on an interface for loopback detection to take effect.

Configuring UDLD Protocol Intervals Use the Administration > UDLD > Configure Global page to configure the UniDirectional Link Detection message probe interval, detection interval, and recovery interval.

Parameters

These parameters are displayed:

- ◆ **Message Interval** – Configures the message interval between UDLD probe messages for ports in the advertisement phase and determined to be bidirectional. (Range: 7-90 seconds; Default: 15 seconds)

UDLD probe messages are sent after linkup or detection phases. During the detection phase, messages are exchanged at the maximum rate of one per second. After that, if the protocol reaches a stable state and determines that the link is bidirectional, the message interval is increased to a configurable value based on a curve known as $M1(t)$, a time-based function described in RFC 5171.

If the link is deemed anything other than bidirectional at the end of the detection phase, this curve becomes a flat line with a fixed value of M_{fast} (7 seconds).

If the link is instead deemed bidirectional, the curve will use M_{fast} for the first four subsequent message transmissions and then transition to an M_{slow} value for all other steady-state transmissions. M_{slow} is the value configured by this command.

- ◆ **Detection Interval** – Sets the amount of time the switch remains in detection state after discovering a neighbor. (Range: 5-255 seconds; Default: 5 seconds)

When a neighbor device is discovered by UDLD, the switch enters “detection state” and remains in this state for specified detection-interval. After the detection-interval expires, the switch tries to decide whether or the link is unidirectional based on the information collected during the “detection state.”

- ◆ **Recovery Status** – Configures the switch to automatically recover from UDLD disabled port state after a period specified by the Recovery Interval. (Default: Disabled)

When automatic recovery state is changed, any ports shut down by UDLD will be reset.

- ◆ **Recovery Interval** – Specifies the period after which to automatically recover from UDLD disabled port state. (Range: 30-86400 seconds; Default: 7 seconds)

When the recovery interval is changed, any ports shut down by UDLD will be reset.

Web Interface

To configure the UDLD message probe interval, detection interval, and recovery interval:

1. Click Administration, UDLD, Configure Global.
2. Select Configure Global from the Step list.
3. Configure the message and detection intervals.
4. Enable automatic recovery if required, and set the recovery interval.
5. Click Apply.

Figure 358: Configuring UDLD Protocol Intervals

Administration > UDLD

Step: 1. Configure Global

Message Interval (7-90)	15	seconds
Detection Interval (5-255)	5	seconds
Recovery Status	☒ Enabled	
Recovery Interval (30-86400)	300	seconds

Apply Revert

Configuring UDLD Interface Settings

Use the Administration > UDLD (Configure Interface) page to enable UDLD and aggressive mode which reduces the shut-down delay after loss of bidirectional connectivity is detected.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-28/52)
- ◆ **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 359: Configuring UDLD Interface Settings



The screenshot shows the 'Administration > UDLD' web interface. At the top, there is a breadcrumb trail and a 'Step: 2. Configure Interface' dropdown. Below this is a 'Port Configuration List' table with a 'Total: 50' count and five pagination buttons (1, 2, 3, 4, 5). The table has seven columns: Port, UDLD, Aggressive Mode, Operation State, Port State, Message Interval (seconds), and Detection Interval (seconds). The first five rows of the table are visible, showing ports 1 through 5, all with UDLD and Aggressive Mode enabled, Operation State set to 'Disabled', Port State set to 'Unknown', Message Interval set to 7 seconds, and Detection Interval set to 5 seconds.

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 360: Displaying UDLD Neighbor Information



Entry	Device ID	Port ID	Device Name	Neighbor State	Expire (seconds)	Message Interval (seconds)	Detection Interval (seconds)
1	S123456	Eth 1/1	ECS4110-52P	Bidirectional	21	7	5

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.
- ◆ General loopback detection provided by the commands described in this section and loopback detection provided by the spanning tree protocol cannot both be enabled at the same time. If loopback detection is enabled for the spanning tree protocol, general loopback detection cannot be enabled on the same interface.
- ◆ When a loopback event is detected on an interface or when an interface is released from a shutdown state caused by a loopback event, a trap message is sent and the event recorded in the system log.
- ◆ Loopback detection must be enabled both globally and on an interface for loopback detection to take effect.

Configuring 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)

When the loopback detection mode is changed (enabled or disabled), any ports placed in shutdown state by the loopback detection process will be immediately restored to operation regardless of the remaining recover time.

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 is 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 361: Configuring Global Settings for LBD

The screenshot shows the 'Administration > LBD' configuration page. At the top, there is a 'Step:' dropdown menu set to '1. Configure Global'. Below this, the 'Global Status' section has an 'Enabled' checkbox. The 'Transmit Interval (1-32767)' is set to 10 seconds. The 'Recover Time (60-1000000)' is checked and set to 60 seconds. The 'Action' dropdown is set to 'Shutdown', and the 'Trap' dropdown is set to 'None'. At the bottom right, there are 'Apply' and 'Revert' 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 which are looped back.

Parameters

These parameters are displayed:


- ◆ **Interface** – Displays a list of ports or trunks.
 - **Port** – Port Identifier. (Range: 1-10/26/28/52)
 - **Trunk** – Trunk Identifier. (Range: 1-8)
- ◆ **Admin State** – Manually enables or disables an interface. (Default: Enabled)
- ◆ **Operation State** – Valid states include Normal or Looped.
- ◆ **Looped VLAN** – Shows the VLANs which are in 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 362: Configuring Interface Settings for LBD



Administration > LBD			
Step: 2. Configure Interface			
Interface ☒ Port ☐ Trunk			
Port List Total: 28			
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

Multicast Filtering


This chapter describes how to configure the following multicast services:

- ◆ [IGMP Snooping](#) – Configures snooping and query parameters.
- ◆ [Filtering and Throttling](#) – Filters specified multicast service, or throttles the maximum of multicast groups allowed on an interface.
- ◆ [MLD Snooping](#) – Configures snooping and query parameters for IPv6.
- ◆ [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 363: 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).

You can also configure a single network-wide multicast VLAN shared by hosts residing in other standard or private VLAN groups, preserving security and data isolation [“Multicast VLAN Registration for IPv4” on page 596](#).

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 562](#)) 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 566](#)). 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 562](#)).

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 566](#)). 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 568](#)).

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: Enabled)

When IGMP snooping is enabled globally, the per VLAN interface settings for IGMP snooping take precedence (see [“Setting IGMP Snooping Status per Interface” on page 570](#)).

When IGMP snooping is disabled globally, snooping can still be configured per VLAN interface, but the interface settings will not take effect until snooping is re-enabled globally.

- ◆ **Proxy Reporting Status** – Enables IGMP Snooping with Proxy Reporting. (Default: Disabled)

When proxy reporting is enabled with this command, the switch performs “IGMP Snooping with Proxy Reporting” (as defined in DSL Forum TR-101, April 2006), including last leave, and query suppression.

Last leave sends out a proxy query when the last member leaves a multicast group, and query suppression means that specific queries are not forwarded from an upstream multicast router to hosts downstream from this device.

When proxy reporting is disabled, all IGMP reports received by the switch are forwarded natively to the upstream multicast routers.

- ◆ **TCN Flood** – Enables flooding of multicast traffic if a spanning tree topology change notification (TCN) occurs. (Default: Disabled)

When a spanning tree topology change occurs, the multicast membership information learned by switch may be out of date. For example, a host linked to one port before the topology change (TC) may be moved to another port after the change. To ensure that multicast data is delivered to all receivers, by default, a switch in a VLAN (with IGMP snooping enabled) that receives a Bridge Protocol Data Unit (BPDU) with TC bit set (by the root bridge) will enter into “multicast flooding mode” for a period of time until the topology has stabilized and the new locations of all multicast receivers are learned.

If a topology change notification (TCN) is received, and all the uplink ports are subsequently deleted, a time out mechanism is used to delete all of the currently learned multicast channels.

When a new uplink port starts up, the switch sends unsolicited reports for all currently learned channels out the new uplink port.

By default, the switch immediately enters into “multicast flooding mode” when a spanning tree topology change occurs. In this mode, multicast traffic will be flooded to all VLAN ports. If many ports have subscribed to different multicast groups, flooding may cause excessive packet loss on the link between the switch and the end host. Flooding may be disabled to avoid this, causing multicast traffic to be delivered only to those ports on which multicast group members have been learned. Otherwise, the time spent in flooding mode can be manually configured to reduce excessive loading.

When the spanning tree topology changes, the root bridge sends a proxy query to quickly re-learn the host membership/port relations for multicast channels. The root bridge also sends an unsolicited Multicast Router Discover (MRD) request to quickly locate the multicast routers in this VLAN.

The proxy query and unsolicited MRD request are flooded to all VLAN ports except for the receiving port when the switch receives such packets.

- ◆ **TCN Query Solicit** – Sends out an IGMP general query solicitation when a spanning tree topology change notification (TCN) occurs. (Default: Disabled)

When the root bridge in a spanning tree receives a TCN for a VLAN where IGMP snooping is enabled, it issues a global IGMP leave message (or query solicitation). When a switch receives this solicitation, it floods it to all ports in the VLAN where the spanning tree change occurred. When an upstream multicast router receives this solicitation, it immediately issues an IGMP general query.

A query solicitation can be sent whenever the switch notices a topology change, even if it is not the root bridge in spanning tree.

- ◆ **Router Alert Option** – Discards any IGMPv2/v3 packets that do not include the Router Alert option. (Default: Disabled)

As described in Section 9.1 of RFC 3376 for IGMP Version 3, the Router Alert Option can be used to protect against DOS attacks. One common method of attack is launched by an intruder who takes over the role of querier, and starts overloading multicast hosts by sending a large number of group-and-source-specific queries, each with a large source list and the Maximum Response Time set to a large value.

To protect against this kind of attack, (1) routers should not forward queries. This is easier to accomplish if the query carries the Router Alert option. (2) Also, when the switch is acting in the role of a multicast host (such as when using proxy routing), it should ignore version 2 or 3 queries that do not contain the Router Alert option.

- ◆ **Unregistered Data Flooding** – Floods unregistered multicast traffic into the attached VLAN. (Default: Disabled)

Once the table used to store multicast entries for IGMP snooping and multicast routing is filled, no new entries are learned. If no router port is configured in the attached VLAN, and unregistered-flooding is disabled, any subsequent multicast traffic not found in the table is dropped, otherwise it is flooded throughout the VLAN.

- ◆ **Forwarding Priority** – Assigns a CoS priority to all multicast traffic. (Range: 0-7, where 7 is the highest priority; Default: Disabled)

This parameter can be used to set a high priority for low-latency multicast traffic such as a video-conference, or to set a low priority for normal multicast traffic not sensitive to latency.

- ◆ **Version Exclusive** – Discards any received IGMP messages which use a version different to that currently configured by the IGMP Version attribute. (Default: Disabled)

- ◆ **IGMP Unsolicited Report Interval** – Specifies how often the upstream interface should transmit unsolicited IGMP reports when proxy reporting is enabled. (Range: 1-65535 seconds, Default: 400 seconds)

When a new upstream interface (that is, uplink port) starts up, the switch sends unsolicited reports for all currently learned multicast channels via the new upstream interface.

This command only applies when proxy reporting is enabled.

- ◆ **Router Port Expire Time** – The time the switch waits after the previous querier stops before it considers it to have expired. (Range: 1-65535, Recommended Range: 300-500 seconds, Default: 300)

- ◆ **IGMP Snooping Version** – Sets the protocol version for compatibility with other devices on the network. This is the IGMP Version the switch uses to send snooping reports. (Range: 1-3; Default: 2)

This attribute configures the IGMP report/query version used by IGMP snooping. Versions 1 - 3 are all supported, and versions 2 and 3 are backward compatible, so the switch can operate with other devices, regardless of the snooping version employed.

- ◆ **Querier Status** – When enabled, the switch can serve as the Querier, which is responsible for asking hosts if they want to receive multicast traffic. This feature is not supported for IGMPv3 snooping. (Default: Disabled)

Web Interface

To configure general settings for IGMP Snooping and Query:

1. Click Multicast, IGMP Snooping, General.
2. Adjust the IGMP settings as required.
3. Click Apply.

Figure 364: Configuring General Settings for IGMP Snooping

Multicast > IGMP Snooping > General	
IGMP Snooping Status	☒ Enabled
Proxy Reporting Status	☐ Enabled
TCN Flood	☐ Enabled
TCN Query Solicit	☐ Enabled
Router Alert Option	☐ Enabled
Unregistered Data Flooding	☐ Enabled
Forwarding Priority (0-7)	☐
Version Exclusive	☐ Enabled
IGMP Unsolicited Report Interval (1-65535)	400 seconds
Router Port Expire Time (1-65535)	300 seconds
IGMP Snooping Version (1-3)	2
Querier Status	☐ Enabled

Apply Revert

Specifying Static Interfaces for a Multicast Router

Use the Multicast > IGMP Snooping > Multicast Router (Add Static Multicast Router) page to statically attach an interface to a multicast router/switch.

Depending on network connections, IGMP snooping may not always be able to locate the IGMP querier. Therefore, if the IGMP querier is a known multicast router/switch connected over the network to an interface (port or trunk) on the switch, the interface (and a specified VLAN) can be manually configured to join all the current multicast groups supported by the attached router. This can ensure that multicast traffic is passed to all the appropriate interfaces within the switch.

Command Usage

IGMP Snooping must be enabled globally on the switch (see [“Configuring IGMP Snooping and Query Parameters” on page 562](#)) 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 365: Configuring a Static Interface for a Multicast Router

The screenshot shows a web interface for configuring 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" field with a dropdown menu set to "1". Below the VLAN field, there is an "Interface" section with two radio buttons: "Port" (which is selected) and "Trunk". The "Port" radio button is accompanied by a dropdown menu set to "1". The "Trunk" radio button is accompanied by a dropdown menu. At the bottom right of the form, there are two buttons: "Apply" and "Revert".

To show the static interfaces attached to a multicast router:

1. Click Multicast, IGMP Snooping, Multicast Router.
2. Select Show Static Multicast Router from the Action list.
3. Select the VLAN for which to display this information.

Figure 366: Showing Static Interfaces Attached a Multicast Router

The screenshot shows the 'Multicast > IGMP Snooping > Multicast Router' configuration page. The 'Action' dropdown is set to 'Show Static Multicast Router'. The 'VLAN' dropdown is set to '1'. Below this, a table titled 'Static Multicast Router Interface List' shows a total of 6 interfaces. The table has two columns: a checkbox and the interface name. The interfaces listed are Unit 1 / Port 1, Unit 1 / Port 2, Unit 1 / Port 3, Trunk 2, Trunk 5, and Unit 1 / Port 4. At the bottom right, there are 'Delete' and 'Revert' buttons.

	Interface
☐	Unit 1 / Port 1
☐	Unit 1 / Port 2
☐	Unit 1 / Port 3
☐	Trunk 2
☐	Trunk 5
☐	Unit 1 / Port 4

Multicast routers that are attached to ports on the switch use information obtained from IGMP, along with a multicast routing protocol (such as PIM) to support IP multicasting across the Internet. These routers may be dynamically discovered by the switch or statically assigned to an interface on the switch.

To show 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 367: Showing Current Interfaces Attached a Multicast Router

The screenshot shows the 'Multicast > IGMP Snooping > Multicast Router' configuration page. The 'Action' dropdown is set to 'Show Current Multicast Router'. The 'VLAN' dropdown is set to '1'. Below this, a table titled 'Multicast Router Interface Information' shows a total of 2 interfaces. The table has three columns: Interface, Type, and Expire. The interfaces listed are Eth 1 / 1 (Static) and Eth 1 / 4 (Dynamic, with an expire time of 0:4:59). A green help icon is visible in the top right corner.

Interface	Type	Expire
Eth 1 / 1	Static	
Eth 1 / 4	Dynamic	0:4:59

Assigning Interfaces to Multicast Services

Use the Multicast > IGMP Snooping > IGMP Member (Add Static Member) page to statically assign a multicast service to an interface.

Multicast filtering can be dynamically configured using IGMP Snooping and IGMP Query messages (see [“Configuring IGMP Snooping and Query Parameters” on page 562](#)). 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 368: Assigning an Interface to a Multicast Service

The screenshot shows a web interface titled "Multicast > IGMP Snooping > IGMP Member". Below the title bar, there is a section labeled "Action:" with a dropdown menu set to "Add Static Member". Below this, there are four fields: "VLAN" with a dropdown menu set to "1", "Interface" with a radio button selected for "Port" and a dropdown menu set to "1", "Multicast IP" with a text input field containing "224.1.1.1", and "Trunk" with a radio button and a dropdown menu set to "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 369: Showing Static Interfaces Assigned to a Multicast Service



	Interface	Multicast IP
☐	Unit 1 / Port 1	224.1.1.1
☐	Unit 1 / Port 2	224.1.2.2
☐	Unit 1 / Port 3	230.1.1.1
☐	Trunk 2	230.1.2.2
☐	Trunk 5	239.1.1.1
☐	Unit 1 / Port 4	239.2.2.2

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 562.

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: Enabled)

When IGMP snooping is enabled globally (see [page 562](#)), 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 562](#)), or as a proxy host when IGMP snooping proxy reporting is enabled ([page 562](#)).

- ◆ **Query Response Interval** – The maximum time the system waits for a response to general queries. (Range: 10-31740 tenths of a second in multiples of 10; Default: 10 seconds)

This attribute applies when the switch is serving as the querier ([page 562](#)), or as a proxy host when IGMP snooping proxy reporting is enabled ([page 562](#)).

- ◆ **Last Member Query Interval** – The interval to wait for a response to a group-specific or group-and-source-specific query message. (Range: 1-31744 tenths of a second in multiples of 10; Default: 1 second)

When a multicast host leaves a group, it sends an IGMP leave message. When the leave message is received by the switch, it checks to see if this host is the last to leave the group by sending out an IGMP group-specific or group-and-source-specific query message, and starts a timer. If no reports are received before the timer expires, the group record is deleted, and a report is sent to the upstream multicast router.

A reduced value will result in reduced time to detect the loss of the last member of a group or source, but may generate more burst traffic.

This attribute will take effect only if IGMP snooping proxy reporting is enabled ([page 562](#)) or IGMP querier is enabled ([page 562](#)).

- ◆ **Last Member Query Count** – The number of IGMP proxy group-specific or group-and-source-specific query messages that are sent out before the system assumes there are no more local members. (Range: 1-255; Default: 2)

This attribute will take effect only if IGMP snooping proxy reporting or IGMP querier is enabled.

- ◆ **Proxy Query Address** – A static source address for locally generated query and report messages used by IGMP Proxy Reporting. (Range: Any valid IP unicast address; Default: 0.0.0.0)

IGMP Snooping uses a null IP address of 0.0.0.0 for the source of IGMP query messages which are proxied to downstream hosts to indicate that it is not the elected querier, but is only proxying these messages as defined in RFC 4541. The switch also uses a null address in IGMP reports sent to upstream ports.

Many hosts do not implement RFC 4541, and therefore do not understand query messages with the source address of 0.0.0.0. These hosts will therefore not reply to the queries, causing the multicast router to stop sending traffic to them.

To resolve this problem, the source address in proxied IGMP query messages can be replaced with any valid unicast address (other than the router's own address).

Web Interface

To configure IGMP snooping on a VLAN:

1. Click Multicast, IGMP Snooping, Interface.
2. Select Configure VLAN from the Action list.
3. Select the VLAN to configure and update the required parameters.
4. Click Apply.

Figure 370: Configuring IGMP Snooping on a VLAN

Multicast > IGMP Snooping > Interface

Action: Configure VLAN

VLAN: 1

IGMP Snooping Status: ☒ Enabled

Version Exclusive: Using Global Status

Immediate Leave Status: ☐ Enabled By-Group

Multicast Router Discovery: ☐ Enabled

General Query Suppression: ☐ Enabled

Proxy Reporting: Using Global Status

Interface Version: Using Global Version

Query Interval (2-31744): 125 seconds

Query Response Interval (10-31740): 100 (1/10 seconds, multiple of 10)

Last Member Query Interval (1-31744): 10 (1/10 seconds, multiple of 10)

Last Member Query Count (1-255): 2


Proxy (Query) Address: 0.0.0.0

Apply Revert

To show the interface settings for IGMP snooping:

1. Click Multicast, IGMP Snooping, Interface.
2. Select Show VLAN Information from the Action list.

Figure 371: Showing Interface Settings for IGMP Snooping



The screenshot shows the 'Multicast > IGMP Snooping > Interface' configuration page. At the top, there is a breadcrumb trail and a title bar. Below the title bar, there is an 'Action:' dropdown menu set to 'Show VLAN Information'. The main content area displays a table titled 'IGMP Snooping VLAN List' with a 'Total: 2' indicator. The table has 13 columns: VLAN, IGMP Snooping Status, Immediate Leave Status, Query Interval, Query Response Interval, Last Member Query Interval, Last Member Query Count, Proxy (Query) Address, Proxy Reporting, Multicast Router Discovery, General Query Suppression, Version Exclusive, and Interface Version. There are two rows of data, one for VLAN 1 and one for VLAN 4093. Both rows show 'Enabled' for IGMP Snooping Status, 'Disabled' for Immediate Leave Status, '125' for Query Interval, '100' for Query Response Interval, '10' for Last Member Query Interval, '2' for Last Member Query Count, '0.0.0.0' for Proxy (Query) Address, 'Using global status (Disabled)' for Proxy Reporting, 'Disabled' for Multicast Router Discovery, 'Disabled' for General Query Suppression, 'Using global status (Disabled)' for Version Exclusive, and 'Using global version (2)' for Interface Version.

VLAN	IGMP Snooping Status	Immediate Leave Status	Query Interval	Query Response Interval	Last Member Query Interval	Last Member Query Count	Proxy (Query) Address	Proxy Reporting	Multicast Router Discovery	General Query Suppression	Version Exclusive	Interface Version
1	Enabled	Disabled	125	100	10	2	0.0.0.0	Using global status (Disabled)	Disabled	Disabled	Using global status (Disabled)	Using global version (2)
4093	Enabled	Disabled	125	100	10	2	0.0.0.0	Using global status (Disabled)	Disabled	Disabled	Using global status (Disabled)	Using global version (2)

Filtering IGMP Query Packets and Multicast Data

Use the Multicast > IGMP Snooping > Interface (Configure Interface) page to configure an interface to drop IGMP query packets or multicast data packets.

Parameters

These parameters are displayed:

- ◆ **Interface** – Port or Trunk identifier.
- ◆ **IGMP Query Drop** – Configures an interface to drop any IGMP query packets received on the specified interface. If this switch is acting as a Querier, this prevents it from being affected by messages received from another Querier.
- ◆ **Multicast Data Drop** – Configures an interface to stop multicast services from being forwarded to users attached to the downstream port (i.e., the interfaces specified by this command).

Web Interface

To drop IGMP query packets or multicast data packets:

1. Click Multicast, IGMP Snooping, Interface.
2. Select Configure Interface from the Action list.
3. Select Port or Trunk interface.
4. Enable the required drop functions for any interface.
5. Click Apply.

Figure 372: Dropping IGMP Query or Multicast Data Packets

The screenshot shows the 'Multicast > IGMP Snooping > Interface' configuration page. At the top, there is a breadcrumb trail and a title bar. Below the title bar, there is an 'Action:' dropdown menu set to 'Configure Interface'. Underneath, there are radio buttons for 'Interface', 'Port', and 'Trunk', with 'Interface' selected. A 'Port List' section shows 'Total: 28' and three tabs labeled 1, 2, and 3. The main table has three columns: 'Port', 'IGMP Query Drop', and 'Multicast Data Drop'. Each row represents a port from 1 to 5, and each has a checkbox for 'Enabled' under both the 'IGMP Query Drop' and 'Multicast Data Drop' headers.

Port	IGMP Query Drop	Multicast Data Drop
1	☐ Enabled	☐ Enabled
2	☐ Enabled	☐ Enabled
3	☐ Enabled	☐ Enabled
4	☐ Enabled	☐ Enabled
5	☐ Enabled	☐ Enabled

Displaying Multicast Groups Discovered by IGMP Snooping

Use the Multicast > IGMP Snooping > Forwarding Entry page to display the forwarding entries learned through IGMP Snooping.

Command Usage

To display information about multicast groups, IGMP Snooping must first be enabled on the switch (see [page 562](#)).

Parameters

These parameters are displayed:

- ◆ **VLAN** – An interface on the switch that is forwarding traffic to downstream ports for the specified multicast group address.
- ◆ **Group Address** – IP multicast group address with subscribers directly attached or downstream from the switch, or a static multicast group assigned to this interface.
- ◆ **Interface** – A downstream port or trunk that is receiving traffic for the specified multicast group. This field may include both dynamically and statically configured multicast router ports.
- ◆ **Up Time** – Time that this multicast group has been known.
- ◆ **Expire** – Time until this entry expires.
- ◆ **Count** – The number of times this address has been learned by IGMP snooping.

Web Interface

To show multicast groups learned through IGMP snooping:

1. Click Multicast, IGMP Snooping, Forwarding Entry.
2. Select the VLAN for which to display this information.

Figure 373: Showing Multicast Groups Learned by IGMP Snooping

Multicast > IGMP Snooping > Forwarding Entry

IGMP Snooping Forwarding Entry List Total: 10

VLAN	Group Address	Source Address	Interface	Up Time	Expire	Count
1	224.1.1.1	*	Eth 1 / 9 (Router Port)	00:00:06:46		2 (Port)
			Eth 1 / 11 (Member Port)	00:00:06:46	03:46	1 (Host)
1	224.1.1.2	192.168.1.2	Eth 1 / 9 (Router Port)		02:24	1 (Port)
2	224.1.1.3	*	Eth 1 / 9 (Router Port)	00:00:16:14		1 (Port)
2	239.255.255.250	*	Eth 1 / 9 (Router Port)	00:00:08:47		2 (Port)
			Eth 1 / 11 (Member Port)	00:00:08:47	03:46	1 (Host)

Click this button to clear all IGMP Snooping dynamic groups.

Displaying IGMP Snooping Statistics Use the Multicast > IGMP Snooping > Statistics pages to display IGMP snooping protocol-related statistics for the specified interface.

Parameters

These parameters are displayed:

- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Port** – Port identifier. (Range: 1-28/52)
- ◆ **Trunk** – Trunk identifier. (Range: 1-26)

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 374: Displaying IGMP Snooping Statistics – Query

The screenshot displays the 'Multicast > IGMP Snooping > Statistics' web interface. At the top, the breadcrumb path is shown. Below it, the 'Action:' dropdown is set to 'Show Query Statistics'. A 'VLAN' dropdown is set to '1'. The main section, titled 'Query Statistics', contains a table of statistics. The table lists various metrics such as 'Other Querier', 'Self Querier', 'General Query Received', and 'Warn Rate Limit' with their corresponding values. At the bottom, there is a 'Clear All' button with a tooltip and a 'Refresh' button.

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 375: Displaying IGMP Snooping Statistics – VLAN

The screenshot shows a web interface for displaying IGMP Snooping Statistics for a specific VLAN. The breadcrumb trail at the top is "Multicast > IGMP Snooping > Statistics". Below this, there is an "Action:" dropdown menu set to "Show VLAN Statistics". A "VLAN" dropdown menu is set to "1". The statistics are organized into two sections: "Input Statistics" and "Output Statistics". Each section contains a table of metrics and their values.

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

At the bottom right of the statistics area, there are two buttons: "Clear" and "Refresh".

To display IGMP snooping protocol-related statistics for a port:

1. Click Multicast, IGMP Snooping, Statistics.
2. Select Show Port Statistics from the Action list.
3. Select a Port.

Figure 376: Displaying IGMP Snooping Statistics – Port

The screenshot shows a web interface for displaying IGMP Snooping statistics. The breadcrumb trail is "Multicast > IGMP Snooping > Statistics". Below this, there is an "Action:" dropdown menu set to "Show Port Statistics". A "Port" 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 with values of 0.

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

At the bottom right of the statistics area is a "Refresh" button.

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 377: Enabling IGMP Filtering and Throttling

The screenshot shows a web interface for configuring IGMP filtering. The breadcrumb trail at the top is "Multicast > IGMP Snooping > Filter". Below this, there is a "Step:" dropdown menu set to "1. Configure General". The main configuration area shows "IGMP Filter Status" with a checked checkbox and the word "Enabled". At the bottom right, there are two buttons: "Apply" and "Revert".

Configuring IGMP Filter Profiles

Use the Multicast > IGMP Snooping > Filter (Configure Profile – Add) page to create an IGMP profile and set its access mode. Then use the (Add Multicast Group Range) page to configure the multicast groups to filter.

Command Usage

Specify a range of multicast groups by entering a start and end IP address; or specify a single multicast group by entering the same IP address for the start and end of the range.

Parameters

These parameters are displayed:

Add

- ◆ **Profile ID** – Creates an IGMP profile. (Range: 1-4294967295)
- ◆ **Access Mode** – Sets the access mode of the profile; either permit or deny. (Default: Deny)

When the access mode is set to permit, IGMP join reports are processed when a multicast group falls within the controlled range. When the access mode is set to deny, IGMP join reports are only processed when the multicast group is not in the controlled range.

Add Multicast Group Range

- ◆ **Profile ID** – Selects an IGMP profile to configure.
- ◆ **Start Multicast IP Address** – Specifies the starting address of a range of multicast groups.
- ◆ **End Multicast IP Address** – Specifies the ending address of a range of multicast groups.

Web Interface

To create an IGMP filter profile and set its access mode:

1. Click Multicast, IGMP Snooping, Filter.
2. Select Configure Profile from the Step list.
3. Select Add from the Action list.
4. Enter the number for a profile, and set its access mode.
5. Click Apply.

Figure 378: Creating an IGMP Filtering Profile

The screenshot shows a web interface for configuring IGMP filtering. The breadcrumb navigation 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 379: Showing the IGMP Filtering Profiles Created

Multicast > IGMP Snooping > Filter

Step: 2. Configure Profile Action: Show

IGMP Snooping Filter Profile List Total: 1

☐	Profile ID	Action Mode
☐	19	Permit

Delete Revert

To add a range of multicast groups to an IGMP filter profile:

1. Click Multicast, IGMP Snooping, Filter.
2. Select Configure Profile from the Step list.
3. Select Add Multicast Group Range from the Action list.
4. Select the profile to configure, and add a multicast group address or range of addresses.
5. Click Apply.

Figure 380: Adding Multicast Groups to an IGMP Filtering Profile

Multicast > IGMP Snooping > Filter

Step: 2. Configure Profile Action: Add Multicast Group Range

Profile ID 19

Start Multicast IP Address 239.2.3.1

End Multicast IP Address 239.2.3.200

Apply Revert

To show the multicast groups configured for an IGMP filter profile:

1. Click Multicast, IGMP Snooping, Filter.
2. Select Configure Profile from the Step list.
3. Select Show Multicast Group Range from the Action list.
4. Select the profile for which to display this information.

Figure 381: Showing the Groups Assigned to an IGMP Filtering Profile

The screenshot shows the configuration page for Multicast > IGMP Snooping > Filter. The 'Step' is set to '2. Configure Profile' and the 'Action' is 'Show Multicast Group Range'. The 'Profile ID' is set to '19'. Below this, there is a table titled 'Multicast IP Address Range List' with a 'Total: 1' count. The table has two columns: 'Start Multicast IP Address' and 'End Multicast IP Address'. A single row is displayed with the start address '239.2.3.1' and the end address '239.2.3.200'. There are checkboxes in the first column of the table. At the bottom of the table, there are 'Delete' and 'Revert' buttons.

	Start Multicast IP Address	End Multicast IP Address
☐	239.2.3.1	239.2.3.200

Configuring IGMP Filtering and Throttling for Interfaces

Use the Multicast > IGMP Snooping > Filter (Configure Interface) page to assign and IGMP filter profile to interfaces on the switch, or to throttle multicast traffic by limiting the maximum number of multicast groups an interface can join at the same time.

Command Usage

- ◆ IGMP throttling sets a maximum number of multicast groups that a port can join at the same time. When the maximum number of groups is reached on a port, the switch can take one of two actions; either “deny” or “replace.” If the action is set to deny, any new IGMP join reports will be dropped. If the action is set to replace, the switch randomly removes an existing group and replaces it with the new multicast group.

Parameters

These parameters are displayed:

- ◆ **Interface** – Port or trunk identifier.
An IGMP profile or throttling setting can be applied to a port or trunk. When ports are configured as trunk members, the trunk uses the settings applied to the first port member in the trunk.
- ◆ **Profile ID** – Selects an existing profile to assign to an interface.
- ◆ **Max Multicast Groups** – Sets the maximum number of multicast groups an interface can join at the same time. (Range: 1-255; 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 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 382: Configuring IGMP Filtering and Throttling Interface Settings

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

MLD Snooping (Snooping and Query for IPv6)

Multicast Listener Discovery (MLD) snooping operates on IPv6 traffic and performs a similar function to IGMP snooping for IPv4. That is, MLD snooping dynamically configures switch ports to limit IPv6 multicast traffic so that it is forwarded only to ports with users that want to receive it. This reduces the flooding of IPv6 multicast packets in the specified VLANs.

There are two versions of the MLD protocol, version 1 and version 2. MLDv1 control packets include Listener Query, Listener Report, and Listener Done messages (equivalent to IGMPv2 query, report, and leave messages). MLDv2 control packets

include MLDv2 query and report messages, as well as MLDv1 report and done messages.

Remember that IGMP Snooping and MLD Snooping are independent functions, and can therefore both function at the same time.

Configuring MLD Snooping and Query Parameters

Use the Multicast > MLD Snooping > General page to configure the switch to forward multicast traffic intelligently. Based on the MLD query and report messages, the switch forwards multicast traffic only to the ports that request it. This prevents the switch from broadcasting the traffic to all ports and possibly disrupting network performance.

Parameters

These parameters are displayed:

- ◆ **MLD Snooping Status** – When enabled, the switch will monitor network traffic to determine which hosts want to receive multicast traffic. (Default: Disabled)
- ◆ **Querier Status** – When enabled, the switch can serve as the querier for MLDv2 snooping if elected. The querier is responsible for asking hosts if they want to receive multicast traffic. (Default: Disabled)

An IPv6 address must be configured on the VLAN interface from which the querier will act if elected. When serving as the querier, the switch uses this IPv6 address as the query source address.

The querier will not start or will disable itself after having started if it detects an IPv6 multicast router on the network.

- ◆ **Robustness** – MLD Snooping robustness variable. A port will be removed from the receiver list for a multicast service when no MLD reports are detected in response to a number of MLD queries. The robustness variable sets the number of queries on ports for which there is no report. (Range: 2-10 Default: 2)

- ◆ **Query Interval** – The interval between sending MLD general queries. (Range: 60-125 seconds; Default: 125 seconds)

This attribute applies when the switch is serving as the querier.

An MLD general query message is sent by the switch at the interval specified by this attribute. When this message is received by downstream hosts, all receivers build an MLD report for the multicast groups they have joined.

- ◆ **Query Max Response Time** – The maximum response time advertised in MLD general queries. (Range: 5-25 seconds; Default: 10 seconds)

This attribute controls how long the host has to respond to an MLD Query message before the switch deletes the group if it is the last member.

- ◆ **Router Port Expiry Time** – The time the switch waits after the previous querier stops before it considers the router port (i.e., the interface that had been

receiving query packets) to have expired. (Range: 300-500 seconds;
Default: 300 seconds)

- ◆ **MLD Snooping Version** – The protocol version used for compatibility with other devices on the network. This is the MLD version the switch uses to send snooping reports. (Range: 1-2; Default: 2)
- ◆ **Unknown Multicast Mode** – The action for dealing with unknown multicast packets. Options include:
 - **Flood** – Floods any received IPv6 multicast packets that have not been requested by a host to all ports in the VLAN.
 - **To Router Port** – Forwards any received IPv6 multicast packets that have not been requested by a host to ports that are connected to a detected multicast router. (This is the default action.)

Web Interface

To configure general settings for MLD Snooping:

1. Click Multicast, MLD Snooping, General.
2. Adjust the settings as required.
3. Click Apply.

Figure 383: 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

Apply Revert

Setting Immediate Leave Status for MLD Snooping per Interface

Use the Multicast > MLD Snooping > Interface page to configure Immediate Leave status for a VLAN.

Parameters

These parameters are displayed:

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

- ◆ **Immediate Leave Status** – Immediately deletes a member port of an IPv6 multicast service when a leave packet is received at that port and immediate leave is enabled for the parent VLAN. (Default: Disabled)

If MLD immediate-leave is *not* used, a multicast router (or querier) will send a group-specific query message when an MLD group leave message is received. The router/querier stops forwarding traffic for that group only if no host replies to the query within the specified timeout period.

If MLD immediate-leave is enabled, the switch assumes that only one host is connected to the interface. Therefore, immediate leave should only be enabled on an interface if it is connected to only one MLD-enabled device, either a service host or a neighbor running MLD snooping.

Web Interface

To configure immediate leave for MLD Snooping:

1. Click Multicast, MLD Snooping, Interface.
2. Select a VLAN, and set the status for immediate leave.
3. Click Apply.

Figure 384: Configuring Immediate Leave for MLD Snooping

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 588](#)) 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 385: 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 set to '1'. The 'Interface' section has two radio buttons: 'Port' (selected) and 'Trunk'. The 'Port' radio button is followed by a dropdown set to '1'. The 'Trunk' radio button is followed by a dropdown. 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 386: 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 set to '1'. The page displays a table titled 'Static Multicast Router Interface List' with a 'Total: 2' count. The table has two columns: a checkbox column and an 'Interface' column. The table contains two rows of data. 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 387: Showing Current Interfaces Attached an IPv6 Multicast Router

Multicast > MLD Snooping > Multicast Router	
Action: Show Current Multicast Router	
VLAN: 1	
Multicast Router Interface Information Total: 4	
Interface	Type
Unit 1 / Port 4	Static
Unit 1 / Port 5	Dynamic
Trunk 2	Dynamic
Trunk 3	Dynamic

Assigning Interfaces to IPv6 Multicast Services

Use the Multicast > MLD Snooping > MLD Member (Add Static Member) page to statically assign an IPv6 multicast service to an interface.

Multicast filtering can be dynamically configured using MLD snooping and query messages (see [“Configuring MLD Snooping and Query Parameters” on page 588](#)). 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 388: 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 'FF00:0:0:0:0:10C'. The 'Interface' is set to 'Port 1' (selected) and 'Trunk' (deselected). There are 'Apply' and 'Revert' 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 389: Showing Static Interfaces Assigned to an IPv6 Multicast Service

The screenshot shows the 'Multicast > MLD Snooping > MLD Member' configuration page. The 'Action' dropdown is set to 'Show Static Member'. The 'VLAN' is set to '1'. Below the form, there is a table titled 'MLD Member Interface List' with a total of 8 entries. The table has three columns: a checkbox, 'Multicast IPv6 Address', and 'Interface'. At the bottom, there are 'Delete' and 'Revert' buttons.

	Multicast IPv6 Address	Interface
☐	FF02::01:01:01:01	Unit 1 / Port 1
☐	FF02::01:01:01:02	Unit 1 / Port 2
☐	FF01::1	Unit 1 / Port 12
☐	FF01::2	Unit 1 / Port 13
☐	FF01::3	Unit 1 / Port 14
☐	FF01::4	Unit 1 / Port 15
☐	FF01::5	Unit 1 / Port 16
☐	FF02::01:01:01:FF	Trunk 3

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 390: 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	MLD Snooping
FF01::1	Trunk 5	MLD Snooping

Showing MLD Snooping Groups and Source List

Use the Multicast > MLD Snooping > Group Information page to display known multicast groups, member ports, the means by which each group was learned, and the corresponding source list.

Parameters

These parameters are displayed:

- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Interface** – Port or trunk identifier.
- ◆ **Group Address** – The IP address for a specific multicast service.
- ◆ **Type** – The means by which each group was learned – MLD Snooping or Multicast Data.
- ◆ **Filter Mode** – The filter mode is used to summarize the total listening state of a multicast address to a minimum set such that all nodes' listening states are respected. In Include mode, the router only uses the request list, indicating that the reception of packets sent to the specified multicast address is requested only from those IP source addresses listed in the hosts' source-list. In Exclude mode, the router uses both the request list and exclude list, indicating that the

reception of packets sent to the given multicast address is requested from all IP source addresses, except for those listed in the exclude source-list and for any other sources where the source timer status has expired.

- ◆ **Filter Timer Elapse** – The Filter timer is only used when a specific multicast address is in Exclude mode. It represents the time for the multicast address filter mode to expire and change to Include mode.
- ◆ **Request List** – Sources included on the router's request list.
- ◆ **Exclude List** – Sources included on the router's exclude list.

Web Interface

To display known MLD multicast groups:

1. Click Multicast, MLD Snooping, Group Information.
2. Select the port or trunk, and then select a multicast service assigned to that interface.

Figure 391: Showing IPv6 Multicast Services and Corresponding Sources

Multicast > MLD Snooping > Group Information

VLAN: 1

Interface: ☒ Port 1 ☐ Trunk 1

Group Address: FF02::01:01:01:01

Type: MLD Snooping

Filter Mode: Exclude

Filter Timer: 10 seconds

Elapse:

Request List Total: 3

IPv6 Address
::01:02:03:01
::01:02:03:02
::01:02:03:03

Exclude List Total: 3


IPv6 Address
::02:02:03:01
::02:02:03:02
::02:02:03:03

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 392: 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 599](#)).
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 600](#)).
3. Set the interfaces that will join the MVR as source ports or receiver ports (see [“Configuring MVR Interface Status” on page 603](#)).
4. For multicast streams that will run for a long term and be associated with a stable set of hosts, you can statically bind the multicast group to the

participating interfaces (see [“Assigning Static MVR Multicast Groups to Interfaces” on page 605](#)).

- ◆ 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 393: Configuring Global Settings for MVR

The screenshot shows the 'Multicast > MVR' configuration page. At the top, there is a breadcrumb 'Multicast > MVR' and a 'Step:' dropdown menu set to '1. Configure Global'. Below this, the configuration options are as follows:

- Proxy Switching:** A checkbox labeled 'Enabled' is checked.
- Robustness Value (1-255):** A text input field containing the value '1'.
- Proxy Query Interval (2-31744):** A text input field containing the value '125', followed by the unit 'sec'.
- Source Port Mode:** A dropdown menu with 'Always Forward' selected.

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

Configuring MVR Domain Settings

Use the Multicast > MVR (Configure Domain) page to enable MVR globally on the switch, and select the VLAN that will serve as the sole channel for common multicast streams supported by the service provider.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **MVR Status** – When MVR is enabled on the switch, any multicast data associated with an MVR group is sent from all designated source ports, to all receiver ports that have registered to receive data from that multicast group. (Default: Disabled)
- ◆ **MVR VLAN** – Identifier of the VLAN that serves as the channel for streaming multicast services using MVR. MVR source ports should be configured as members of the MVR VLAN (see [“Adding Static Members to VLANs” on page 165](#)), but MVR receiver ports should not be manually configured as members of this VLAN. (Default: 1)
- ◆ **MVR Running Status** – Indicates whether or not all necessary conditions in the MVR environment are satisfied. Running status is Active as long as MVR is enabled, the specified MVR VLAN exists, and a source port with a valid link has been configured (see [“Configuring MVR Interface Status” on page 603](#)).
- ◆ **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 394: Configuring Domain Settings for MVR

The screenshot shows the 'Multicast > MVR' configuration page. The 'Step' dropdown is set to '2. Configure Domain'. The configuration fields are as follows:

Field	Value
Domain ID	1
MVR Status	☒ Enabled
MVR VLAN	1
MVR Running Status	Active
MVR Current Learned Groups	0
Forwarding Priority (0-7)	☐ 0
Upstream Source IP	192.168.0.12

Buttons: Apply, Revert

Configuring MVR Group Address Profiles

Use the Multicast > MVR (Configure Profile and Associate Profile) pages to assign the multicast group address for required services to one or more MVR domains.

Command Usage

- ◆ Use the Configure Profile page to statically configure all multicast group addresses that will join the MVR VLAN. Any multicast data associated with an MVR group is sent from all source ports to all receiver ports that have registered to receive data from that multicast group.
- ◆ The IP address range from 224.0.0.0 to 239.255.255.255 is used for multicast streams. MVR group addresses cannot fall within the reserved IP multicast address range of 224.0.0.x.

- ◆ IGMP snooping and MVR share a maximum number of 1024 groups. Any multicast streams received in excess of this limitation will be flooded to all ports in the associated domain.

Parameters

These parameters are displayed:

Configure Profile

- ◆ **Profile Name** – The name of a profile containing one or more MVR group addresses. (Range: 1-21 characters)
- ◆ **Start IP Address** – Starting IP address for an MVR multicast group. (Range: 224.0.1.0 - 239.255.255.255)
- ◆ **End IP Address** – Ending IP address for an MVR multicast group. (Range: 224.0.1.0 - 239.255.255.255)

Associate Profile

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **Profile Name** – The name of a profile to be assigned to this domain. (Range: 1-21 characters)

Web Interface

To configure an MVR group address profile:

1. Click Multicast, MVR.
2. Select Configure Profile from the Step list.
3. Select Add from the Action list.
4. Enter the name of a group profile to be assigned to one or more domains, and specify a multicast group that will stream traffic to participating hosts.
5. Click Apply.

Figure 395: Configuring an MVR Group Address Profile

The screenshot shows a web interface for configuring an MVR group address profile. The breadcrumb is 'Multicast > MVR'. The 'Step' dropdown is set to '2. Configure Profile' and the 'Action' dropdown is set to 'Add'. The form contains three input fields: 'Profile Name' with the value 'sales', 'Start IP Address' with the value '234.5.6.8', and 'End IP Address' with the value '234.5.6.10'. At the bottom right, there are two buttons: 'Apply' and '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 396: Displaying MVR Group Address Profiles

The screenshot shows the 'Multicast > MVR' configuration page. At the top, there is a breadcrumb 'Multicast > MVR' and a help icon. Below this, there are two dropdown menus: 'Step: 3. Configure Profile' and 'Action: Show'. The main content area is titled 'MVR Profile List Total: 2'. It contains a table with three columns: a checkbox, 'Profile Name', and 'Start IP Address - End IP Address'. There are two rows of data: 'profile1' with IP range '224.1.1.1 - 224.1.1.10' and 'profile2' with IP range '224.1.2.1 - 224.1.2.100'. At the bottom right of the table, there are 'Delete' and 'Revert' buttons.

☐	Profile Name	Start IP Address - End IP Address
☐	profile1	224.1.1.1 - 224.1.1.10
☐	profile2	224.1.2.1 - 224.1.2.100

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 397: Assigning an MVR Group Address Profile to a Domain

The screenshot shows the 'Multicast > MVR' configuration page. At the top, there is a breadcrumb 'Multicast > MVR'. Below this, there are two dropdown menus: 'Step: 4. Associate Profile' and 'Action: Add'. The main content area has two dropdown menus: 'Domain ID' with the value '1' and 'Profile Name' with the value 'sales'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 398: Showing the MVR Group Address Profiles Assigned to a Domain

The screenshot shows the 'Multicast > MVR' configuration page. At the top, there's a breadcrumb 'Multicast > MVR' and a help icon. Below that, the 'Step' is set to '3. Associate Profile' and the 'Action' is 'Show'. A 'Domain ID' dropdown is set to '1'. Below this, a section titled 'Domain Associated Profile List' shows 'Total: 2'. A table lists two profiles: 'rd' with IP range '234.5.6.7' and 'sales' with IP range '234.5.6.8'. At the bottom, there are 'Delete' and 'Revert' buttons.

Profile Name	Start IP Address	End IP Address
rd	234.5.6.7	234.5.6.7
sales	234.5.6.8	234.5.6.10

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 605](#)).

Receiver ports should not be statically configured as a member of the MVR VLAN. If so configured, its MVR status will be inactive. Also, note that VLAN membership for MVR receiver ports cannot be set to access mode (see [“Adding Static Members to VLANs” on page 165](#)).
- ◆ One or more interfaces may be configured as MVR source ports. A source port is able to both receive and send data for configured MVR groups or for groups which have been statically assigned (see [“Assigning Static MVR Multicast Groups to Interfaces” on page 605](#)).

All source ports must belong to the MVR VLAN.
Subscribers should not be directly connected to source ports.
- ◆ Immediate leave applies only to receiver ports. When enabled, the receiver port is immediately removed from the multicast group identified in the leave message. When immediate leave is disabled, the switch follows the standard rules by sending a query message to the receiver port and waiting for a

response to determine if there are any remaining subscribers for that multicast group before removing the port from the group list.

- Using immediate leave can speed up leave latency, but should only be enabled on a port attached to one multicast subscriber to avoid disrupting services to other group members attached to the same interface.
- Immediate leave does not apply to multicast groups which have been statically assigned to a port.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **Port/Trunk** – Interface identifier.
- ◆ **Type** – The following interface types are supported:
 - **Source** – An uplink port that can send and receive multicast data for the groups assigned to the MVR VLAN. Note that the source port must be manually configured as a member of the MVR VLAN (see [“Adding Static Members to VLANs” on page 165](#)).
 - **Receiver** – A subscriber port that can receive multicast data sent through the MVR VLAN. Any port configured as a receiver port will be dynamically added to the MVR VLAN when it forwards an IGMP report or join message from an attached host requesting any of the designated multicast services supported by the MVR VLAN. Just remember that only IGMP version 2 or 3 hosts can issue multicast join or leave messages. If MVR must be configured for an IGMP version 1 host, the multicast groups must be statically assigned (see [“Assigning Static MVR Multicast Groups to Interfaces” on page 605](#)).
 - **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 399: Configuring Interface Settings for MVR

Multicast > MVR

Step: 5. Configure Interface

Domain ID: 1

Interface: ☒ Port ☐ Trunk

Port List Total: 28

Port	Type	Forwarding Status	MVR Status	Immediate Leave
1	Source	Forwarding	Inactive	☐ By-Group
2	Receiver	Forwarding	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 400: Assigning Static MVR Groups to an Interface

Multicast > MVR

Step: 5. Configure Static Group Member Action: Add

Domain ID: 1

Interface: Port 2 Trunk

VLAN: 2

Group IP Address: 234.5.6.8

Apply Revert

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 401: Showing the Static MVR Groups Assigned to a Port

Multicast > MVR

Step: 5. Configure Static Group Member Action: Show

Domain ID: 1

Interface: Port 2 Trunk

MVR Static Group Member List Total: 3

	VLAN	Group IP Address
☐	2	234.5.6.8
☐	2	234.5.6.9
☐	2	234.5.6.10

Delete Revert

Displaying MVR Receiver Groups Use the Multicast > MVR (Show Member) page to show the multicast groups either statically or dynamically assigned to the MVR receiver groups on each interface.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **Group IP Address** – Multicast groups assigned to the MVR VLAN.


- ◆ **VLAN** – The VLAN through which the service is received. Note that this may be different from the MVR VLAN if the group address has been statically assigned.
- ◆ **Port** – Shows the interfaces with subscribers for multicast services provided through the MVR VLAN.
- ◆ **Up Time** – Time this service has been forwarded to attached clients.
- ◆ **Expire** – Time before this entry expires if no membership report is received from currently active or new clients.
- ◆ **Count** – The number of multicast services currently being forwarded from the MVR VLAN.
- ◆ **Clear MVR Group** – Clears multicast group information dynamically learned through MVR6. Statically configured multicast addresses are not cleared.

Web Interface

To display the interfaces assigned to the MVR receiver groups:

1. Click Multicast, MVR.
2. Select Show Member from the Step list.
3. Select an MVR domain.

Figure 402: Displaying MVR Receiver Groups



Group IP Address	VLAN	Port	Up Time	Expire	Count
224.1.1.1	2		00:00:30		2 (Port)
	2	Unit 1 / Port 1 (Source)			
	1	Unit 1 / Port 2 (Receiver)	00:01:10	00:00	4 (Host)
224.1.1.2	4		00:00:50		4 (Port)
	4	Unit 1 / Port 3 (Source)			
	5	Unit 1 / Port 4 (Receiver)			
	6	Unit 1 / Port 5 (Source)			
	7	Unit 1 / Port 6 (Receiver)	00:01:10	00:00	1 (Host)

Displaying MVR Statistics Use the Multicast > MVR > Show Statistics pages to display MVR protocol-related statistics for the specified interface.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)

- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Port** – Port identifier. (Range: 1-28/52)
- ◆ **Trunk** – Trunk identifier. (Range: 1-26)

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 403: Displaying MVR Statistics – Query

The screenshot displays the 'Multicast > MVR' configuration page. At the top, there are two dropdown menus: 'Step:' set to '7. Show Statistics' and 'Action:' set to 'Show Query Statistics'. Below these, a 'Domain ID' dropdown is set to '1'. The main section is titled 'Query Statistics' and contains a table with the following data:

Querier IP Address	None
Querier Expire Time	00(h):00(m):00(s)
General Query Received	0
General Query Sent	0
Specific Query Received	0
Specific Query Sent	0
Number of Reports Sent	0
Number of Leaves Sent	0

At the bottom left, there is a 'Clear All' button with a tooltip that reads: 'Click this button to clear all MVR statistics of the domain.' At the bottom right, there is a 'Refresh' button.

To display MVR protocol-related statistics for a VLAN:

1. Click Multicast, MVR.
2. Select Show Statistics from the Step list.
3. Select Show VLAN Statistics from the Action list.
4. Select an MVR domain.
5. Select a VLAN.

Figure 404: Displaying MVR Statistics – VLAN

The screenshot shows the 'Multicast > MVR' 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		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 page, there is a 'Refresh' button.

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 405: Displaying MVR Statistics – Port

Multicast > MVR

Step: 7. Show Statistics Action: Show Port Statistics

Domain ID1

Port1

Input Statistics

Report	0	Drop	0
Leave	0	Join Success	0
G Query	0	Group	0
G(-S)-S Query	0		

Output Statistics

Report	0
Leave	0
G Query	0
G(-S)-S Query	0

Refresh

Multicast VLAN Registration for IPv6

MVR6 functions in a manner similar to that described for MRV (see [“Multicast VLAN Registration for IPv4” on page 596](#)).

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 615](#)).
 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 616](#)).
 3. Set the interfaces that will join the MVR as source ports or receiver ports (see [“Configuring MVR6 Interface Status” on page 619](#)).
 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 621](#)).

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 406: Configuring Global Settings for MVR6

The screenshot shows a web interface for configuring MVR6 settings. At the top, it says "Multicast > MVR6". Below that is a "Step:" dropdown menu with "1. Configure Global" selected. The main configuration area has four rows of settings:

Proxy Switching	☒ Enabled
Robustness Value (1-10)	1
Proxy Query Interval (2-31744)	125 sec
Source Port Mode	Always Forward

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

Configuring MVR6 Domain Settings

Use the Multicast > MVR6 (Configure Domain) page to enable MVR6 globally on the switch, and select the VLAN that will serve as the sole channel for common multicast streams supported by the service provider.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **MVR6 Status** – When MVR6 is enabled on the switch, any multicast data associated with an MVR6 group is sent from all designated source ports, to all receiver ports that have registered to receive data from that multicast group. (Default: Disabled)
- ◆ **MVR6 VLAN** – Identifier of the VLAN that serves as the channel for streaming multicast services using MVR6. MVR6 source ports should be configured as members of the MVR6 VLAN (see [“Adding Static Members to VLANs” on page 165](#)), but MVR6 receiver ports should not be manually configured as members of this VLAN. (Default: 1)
- ◆ **MVR6 Running Status** – Indicates whether or not all necessary conditions in the MVR6 environment are satisfied. Running status is Active as long as MVR6 is enabled, the specified MVR6 VLAN exists, and a source port with a valid link has been configured (see [“Configuring MVR6 Interface Status” on page 619](#)).
- ◆ **MVR6 Current Learned Groups** – The number of MVR6 groups currently assigned to this domain.
- ◆ **Forwarding Priority** – The CoS priority assigned to all multicast traffic forwarded into this domain. (Range: 0-7, where 7 is the highest priority)

This parameter can be used to set a high priority for low-latency multicast traffic such as a video-conference, or to set a low priority for normal multicast traffic not sensitive to latency.
- ◆ **Upstream Source IPv6** – The source IPv6 address assigned to all MVR6 control packets sent upstream on the specified domain. This parameter must be a full IPv6 address including the network prefix and host address bits. By default, all MVR6 reports sent upstream use a null source IP address.

All IPv6 addresses must be according to RFC 2373 “IPv6 Addressing Architecture,” using 8 colon-separated 16-bit hexadecimal values. One double colon may be used in the address to indicate the appropriate number of zeros required to fill the undefined fields. (Note that the IP address ff02::X is reserved.)

Web Interface

To configure settings for an MVR6 domain:

1. Click Multicast, MVR6.
2. Select Configure Domain from the Step list.

3. Select a domain from the scroll-down list.
4. Enable MVR6 for the selected domain, select the MVR6 VLAN, set the forwarding priority to be assigned to all ingress multicast traffic, and set the source IP address for all control packets sent upstream as required.
5. Click Apply.

Figure 407: Configuring Domain Settings for MVR6

The screenshot shows the 'Multicast > MVR6' configuration page. At the top, there is a 'Step:' dropdown menu set to '2. Configure Domain'. Below this, the configuration fields are as follows:

- Domain ID:** A dropdown menu with '1' selected.
- MVR6 Status:** A checkbox labeled 'Enabled' which is checked.
- MVR6 VLAN:** A dropdown menu with '1' selected.
- MVR6 Running Status:** A text label 'Inactive'.
- MVR6 Current Learned Groups:** A text label '0'.
- Forwarding Priority (0-7):** A checkbox followed by a text input field.
- Upstream Source IPv6:** A text input field containing the address '2001:d88:2222:7223::72'.

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

Configuring MVR6 Group Address Profiles

Use the Multicast > MVR6 (Configure Profile and Associate Profile) pages to assign the multicast group address for required services to one or more MVR6 domains.

Command Usage

- ◆ Use the Configure Profile page to statically configure all multicast group addresses that will join the MVR6 VLAN. Any multicast data associated with an MVR6 group is sent from all source ports to all receiver ports that have registered to receive data from that multicast group.
- ◆ All IPv6 addresses must be according to RFC 2373 "IPv6 Addressing Architecture," using 8 colon-separated 16-bit hexadecimal values. One double colon may be used in the address to indicate the appropriate number of zeros required to fill the undefined fields. (Note that the IP address ff02::X is reserved.)
- ◆ The MVR6 group address range assigned to a profile cannot overlap with the group address range of any other profile.
- ◆ MVR6 domains can be associated with more than one MVR6 profile. But since MVR6 domains cannot share the group range, an MVR6 profile can only be associated with one MVR6 domain.

Parameters

These parameters are displayed:

Configure Profile

- ◆ **Profile Name** – The name of a profile containing one or more MVR6 group addresses. (Range: 1-21 characters)
- ◆ **Start IPv6 Address** – Starting IP address for an MVR6 multicast group. This parameter must be a full IPv6 address including the network prefix and host address bits.
- ◆ **End IPv6 Address** – Ending IP address for an MVR6 multicast group. This parameter must be a full IPv6 address including the network prefix and host address bits.

Associate Profile

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **Profile Name** – The name of a profile to be assigned to this domain. (Range: 1-21 characters)

Web Interface

To configure an MVR6 group address profile:

1. Click Multicast, MVR6.
2. Select Configure Profile from the Step list.
3. Select Add from the Action list.
4. Enter the name of a group profile to be assigned to one or more domains, and specify a multicast group that will stream traffic to participating hosts.
5. Click Apply.

Figure 408: Configuring an MVR6 Group Address Profile

Multicast > MVR6

Step: 2. Configure Profile Action: Add

Profile Name: rd

Start IPv6 Address: ff00::1

End IPv6 Address: ff00::9

Apply Revert

To show the configured MVR6 group address profiles:

1. Click Multicast, MVR6.
2. Select Configure Profile from the Step list.
3. Select Show from the Action list.

Figure 409: Displaying MVR6 Group Address Profiles

The screenshot shows the 'Multicast > MVR6' configuration page. At the top, there is a breadcrumb 'Multicast > MVR6'. Below it, a 'Step:' dropdown is set to '3. Configure Profile' and an 'Action:' dropdown is set to 'Show'. The main section is titled 'MVR6 Profile List' with a 'Total: 1' indicator. It contains a table with the following data:

	Profile Name	Start IPv6 Address	End IPv6 Address
☐	rd	FF00::1	FF00::9

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

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 410: Assigning an MVR6 Group Address Profile to a Domain

The screenshot shows the 'Multicast > MVR6' configuration page. At the top, there is a breadcrumb 'Multicast > MVR6'. Below it, a 'Step:' dropdown is set to '4. Associate Profile' and an 'Action:' dropdown is set to 'Add'. The main section contains two dropdown menus: 'Domain ID' with '1' selected and 'Profile Name' with 'rd' selected. At the bottom right are 'Apply' and 'Revert' buttons.

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 411: Showing MVR6 Group Address Profiles Assigned to a Domain

Multicast > MVR6

Step: 3. Associate Profile Action: Show

Domain ID: 1

Domain Associated Profile List Total: 1

☐	Profile Name	Start IPv6 Address	End IPv6 Address
☐	rd	FF00::1	FF00::9

Buttons: Delete, Revert

Configuring MVR6 Interface Status

Use the Multicast > MVR6 (Configure Interface) page to configure each interface that participates in the MVR6 protocol as a source port or receiver port. If you are sure that only one subscriber attached to an interface is receiving multicast services, you can enable the immediate leave function.

Command Usage

- ◆ A port configured as an MVR6 receiver or source port can join or leave multicast groups configured under MVR6. A port which is not configured as an MVR receiver or source port can use MLD snooping to join or leave multicast groups using the standard rules for multicast filtering (see [“MLD Snooping \(Snooping and Query for IPv6\)” on page 587](#)).

- ◆ 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 605](#)).

Receiver ports should not be statically configured as a member of the MVR6 VLAN. If so configured, its MVR6 status will be inactive. Also, note that VLAN membership for MVR6 receiver ports cannot be set to access mode (see [“Adding Static Members to VLANs” on page 165](#)).

- ◆ One or more interfaces may be configured as MVR6 source ports. A source port is able to both receive and send data for configured MVR6 groups or for groups which have been statically assigned (see [“Assigning Static MVR Multicast Groups to Interfaces” on page 605](#)).

All source ports must belong to the MVR6 VLAN.

Subscribers should not be directly connected to source ports.

- ◆ Immediate leave applies only to receiver ports. When enabled, the receiver port is immediately removed from the multicast group identified in the leave message. When immediate leave is disabled, the switch follows the standard rules by sending a group-specific query to the receiver port and waiting for a

response to determine if there are any remaining subscribers for that multicast group before removing the port from the group list.

- Using immediate leave can speed up leave latency, but should only be enabled on a port attached to one multicast subscriber to avoid disrupting services to other group members attached to the same interface.
- Immediate leave does not apply to multicast groups which have been statically assigned to a port.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **Port/Trunk** – Interface identifier.
- ◆ **Type** – The following interface types are supported:
 - **Non-MVR6** – An interface that does not participate in the MVR6 VLAN. (This is the default type.)
 - **Source** – An uplink port that can send and receive multicast data for the groups assigned to the MVR6 VLAN. Note that the source port must be manually configured as a member of the MVR6 VLAN (see [“Adding Static Members to VLANs” on page 165](#)).
 - **Receiver** – A subscriber port that can receive multicast data sent through the MVR6 VLAN. Also, note that VLAN membership for MVR receiver ports cannot be set to access mode (see [“Adding Static Members to VLANs” on page 165](#)).
- ◆ **Forwarding Status** – Shows if multicast traffic is being forwarded or blocked.
- ◆ **MVR6 Status** – Shows the MVR6 status. MVR6 status for source ports is “Active” if MVR6 is globally enabled on the switch. MVR6 status for receiver ports is “Active” only if there are subscribers receiving multicast traffic from one of the MVR6 groups, or a multicast group has been statically assigned to an interface.
- ◆ **Immediate Leave** – Configures the switch to immediately remove an interface from a multicast stream as soon as it receives a leave message for that group. This option only applies to an interface configured as an MVR6 receiver. (Default: Disabled)
 - **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 412: Configuring Interface Settings for MVR6

Multicast > MVR6

Step: 5. Configure Interface

Domain ID: 1

Interface: ☒ Port ☐ Trunk

Port List Total: 28

Port	Type	Forwarding Status	MVR6 Status	Immediate Leave
1	Source	Forwarding	Inactive	☐ By-Group
2	Receiver	Forwarding	Inactive	☒ By-Group
3	Non-MVR6	Discarding	Inactive	☐ By-Group
4	Non-MVR6	Discarding	Inactive	☐ By-Group
5	Non-MVR6	Discarding	Inactive	☐ By-Group

Assigning Static MVR6 Multicast Groups to Interfaces

Use the Multicast > MVR6 (Configure Static Group Member) page to statically bind multicast groups to a port which will receive long-term multicast streams associated with a stable set of hosts.

Command Usage

- ◆ Multicast groups can be statically assigned to a receiver port using this configuration page.
- ◆ All IPv6 addresses must be according to RFC 2373 "IPv6 Addressing Architecture," using 8 colon-separated 16-bit hexadecimal values. One double colon may be used in the address to indicate the appropriate number of zeros required to fill the undefined fields. (Note that the IP address ff02::X is reserved.)
- ◆ The MVR6 VLAN cannot be specified as the receiver VLAN for static bindings.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **Interface** – Port or trunk identifier.
- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Group IPv6 Address** – Defines a multicast service sent to the selected port. Multicast groups must be assigned from the MVR6 group range configured on the Configure General page.

Web Interface

To assign a static MVR6 group to an interface:

1. Click Multicast, MVR6.
2. Select Configure Static Group Member from the Step list.
3. Select Add from the Action list.
4. Select an MVR6 domain.
5. Select a VLAN and interface to receive the multicast stream, and then enter the multicast group address.
6. Click Apply.

Figure 413: Assigning Static MVR6 Groups to a Port

The screenshot shows a web interface for configuring MVR6. The breadcrumb is 'Multicast > MVR6'. 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:

Domain ID	1
Interface	☒ Port 2 ☐ Trunk
VLAN	2
Group IPv6 Address	ff00::1

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

To show the static MVR6 groups assigned to an interface:

1. Click Multicast, MVR6.
2. Select Configure Static Group Member from the Step list.
3. Select Show from the Action list.

4. Select an MVR6 domain.
5. Select the port or trunk for which to display this information.

Figure 414: Showing the Static MVR6 Groups Assigned to a Port

The screenshot shows the 'Multicast > MVR6' configuration page. At the top, there's a breadcrumb 'Multicast > MVR6' and a help icon. Below that, a 'Step:' dropdown is set to '6. Configure Static Group Member' and an 'Action:' dropdown is set to 'Show'. The 'Domain ID' is set to '1'. The 'Interface' section has radio buttons for 'Port' (selected) and 'Trunk', with 'Port' set to '2'. Below this is a table titled 'MVR6 Static Group Member List' with a 'Total: 1' count. The table has two columns: 'VLAN' and 'Group IPv6 Address'. There is one entry with 'VLAN' set to '2' and 'Group IPv6 Address' set to 'FF00::1'. At the bottom right of the table are 'Delete' and 'Revert' buttons.

VLAN	Group IPv6 Address
2	FF00::1

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 415: Displaying MVR6 Receiver Groups

The screenshot shows the 'Multicast > MVR6' web interface. The 'Step' dropdown is set to '6. Show Member'. The 'Domain ID' is set to '1'. Below this, the 'MVR6 Member List' is displayed with a total of 3 members. The table has columns for Group IPv6 Address, VLAN, Port, Up Time, Expire, and Count.

Group IPv6 Address	VLAN	Port	Up Time	Expire	Count
FF00::1	1		00:00:11:35		2 (Port)
	1	Unit 1 / Port 1 (Source)			
	2	Unit 1 / Port 2 (Receiver)			0 (Host)

Displaying MVR6 Statistics

Use the Multicast > MVR6 > Show Statistics pages to display MVR6 protocol-related statistics for the specified interface.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Port** – Port identifier. (Range: 1-28/52)
- ◆ **Trunk** – Trunk identifier. (Range: 1-26)

Query Statistics

- ◆ **Querier IPv6 Address** – The IP address of the querier on this interface.
- ◆ **Querier Expire Time** – The time after which this querier is assumed to have expired.
- ◆ **General Query Received** – The number of general queries received on this interface.
- ◆ **General Query Sent** – The number of general queries sent from this interface.
- ◆ **Specific Query Received** – The number of specific queries received on this interface.
- ◆ **Specific Query Sent** – The number of specific queries sent from this interface.

- ◆ **Number of Reports Sent** – The number of reports sent from this interface.
- ◆ **Number of Leaves Sent** – The number of leaves sent from this interface.

VLAN, Port, and Trunk Statistics

Input Statistics

- ◆ **Report** – The number of MLD membership reports received on this interface.
- ◆ **Leave** – The number of leave messages received on this interface.
- ◆ **G Query** – The number of general query messages received on this interface.
- ◆ **G(-S)-S Query** – The number of group specific or group-and-source specific query messages received on this interface.
- ◆ **Drop** – The number of times a report, leave or query was dropped. Packets may be dropped due to invalid format, rate limiting, packet content not allowed, or MVR6 group report received.
- ◆ **Join Success** – The number of times a multicast group was successfully joined.
- ◆ **Group** – The number of MVR6 groups active on this interface.

Output Statistics

- ◆ **Report** – The number of MLD membership reports sent from this interface.
- ◆ **Leave** – The number of leave messages sent from this interface.
- ◆ **G Query** – The number of general query messages sent from this interface.
- ◆ **G(-S)-S Query** – The number of group specific or group-and-source specific query messages sent from this interface.

Web Interface

To display statistics for MVR6 query-related messages:

1. Click Multicast, MVR6.
2. Select Show Statistics from the Step list.
3. Select Show Query Statistics from the Action list.
4. Select an MVR6 domain.

Figure 416: Displaying MVR6 Statistics – Query

The screenshot displays the 'Multicast > MVR6' web interface. At the top, there are two dropdown menus: 'Step' set to '7. Show Statistics' and 'Action' set to 'Show Query Statistics'. Below these, a 'Domain ID' dropdown is set to '1'. The main section is titled 'Query Statistics' and contains a table of statistics for Domain ID 1.

Querier IPv6 Address	::
Querier Expire Time	00(h):00(m):00(s)
General Query Received	0
General Query Sent	34
Specific Query Received	0
Specific Query Sent	0
Number of Reports Sent	0
Number of Leaves Sent	0

At the bottom left, there is a 'Clear All' button and a link: 'Click this button to clear all MVR6 statistics of the domain.' At the bottom right, there is a 'Refresh' button.

To display MVR6 protocol-related statistics for a VLAN:

1. Click Multicast, MVR6.
2. Select Show Statistics from the Step list.
3. Select Show VLAN Statistics from the Action list.
4. Select an MVR6 domain.
5. Select a VLAN.

Figure 417: 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 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
Done	0
G Query	0
G(-S)-S Query	0
Drop	0
Join Success	0
Group	1

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 418: Displaying MVR6 Statistics – Port

The screenshot shows the 'Multicast > MVR6' configuration page. At the top, there are two dropdown menus: 'Step: 7. Show Statistics' and 'Action: Show Port Statistics'. Below these, there are two more dropdown menus: 'Domain ID: 1' and 'Port: 1'. The main content area is divided into two sections: 'Input Statistics' and 'Output Statistics'. The 'Input Statistics' section shows four rows of data: 'Report' (0), 'Done' (0), 'G Query' (0), and 'G(-S)-S Query' (0). The 'Output Statistics' section shows four rows of data: 'Report' (12), 'Done' (1), 'G Query' (0), and 'G(-S)-S Query' (0). At the bottom right, there are two buttons: 'Clear' and 'Refresh'.

Multicast > MVR6			
Step: 7. Show Statistics		Action: Show Port Statistics	
Domain ID: 1		Port: 1	
Input Statistics			
Report	0	Drop	0
Done	0	Join Success	0
G Query	0	Group	0
G(-S)-S Query	0		
Output Statistics			
Report	12		
Done	1		
G Query	0		
G(-S)-S Query	0		

Clear Refresh

This chapter provides information on network functions including:

- ◆ **Ping** – Sends ping message to another node on the network.
- ◆ **Trace** – Sends ICMP echo request packets to another node on the network.
- ◆ **Address Resolution Protocol** – Describes how to configure ARP aging time, proxy ARP, or static addresses. Also shows how to display dynamic entries in the ARP cache.

Using the Ping Function

Use the IP > General > Ping page to send ICMP echo request packets to another node on the network.

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.

Parameters

These parameters are displayed:

- ◆ **Host Name/IP Address** – IPv4/IPv6 address or alias of the host.
- ◆ **Probe Count** – Number of packets to send. (Range: 1-16)
- ◆ **Data Size** – Number of bytes in a packet. (Range: 32-1472 bytes for IPv4, 0-512 bytes for IPv6)

The actual packet size will be eight bytes larger than the size specified because the switch adds header information.

Web Interface

To ping another device on the network:

1. Click Tool, Ping.
2. Specify the target device and ping parameters.
3. Click Apply.

Figure 419: Pinging a Network Device

The screenshot displays the 'Tool > Ping' configuration page. It includes input fields for 'Host Name/IP Address', 'Probe Count (1-16)' (set to 5), and 'Data Size (IPv4 : 32-512, IPv6 : 0-1500)' (set to 32 bytes). A note explains that for IPv4, data sizes 0-31 are rounded up to 32 bytes, 32-512 are valid, and 513-1500 are rounded up to 512 bytes. Below the fields are 'Apply' and 'Revert' buttons. The 'Result' section shows a successful ping to 192.168.2.99 with 5 packets received (100%) and a response time of 1000 ms.

Tool > Ping

Host Name/IP Address

Probe Count (1-16)

Data Size (IPv4 : 32-512, IPv6 : 0-1500) bytes

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

Result

```
PING to 192.168.2.99, by 5 of 32-byte payload ICMP packets, timeout is 3 seconds

response time: 1000 ms
response time: 0 ms
response time: 0 ms
response time: 0 ms
response time: 0 ms

Ping statistics for 192.168.2.99:
    5 packets transmitted, 5 packets received (100%), 0 packets lost (0%)
Approximate round trip times:
    Minimum = 0 ms, Maximum = 1000 ms, Average = 200 ms
```

Using the Trace Route Function

Use the IP > General > Trace Route page to show the route packets take to the specified destination.

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.

Parameters

These parameters are displayed:


- ◆ **Destination IP Address** – IPv4/IPv6 address of the host.
- ◆ **IPv4 Max Failures** – The maximum number of failures before which the trace route is terminated. (Fixed: 5)
- ◆ **IPv6 Max Failures** – The maximum number of failures before which the trace route is terminated. (Range: 1-255; Default: 5)

Web Interface

To trace the route to another device on the network:

1. Click IP, General, Trace Route.
2. Specify the target device.
3. Click Apply.

Figure 420: Tracing the Route to a Network Device



Address Resolution Protocol

If IP routing is enabled (page 589), 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 42: 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 421: 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.

Web Interface

To configure the timeout for the ARP cache or to enable Proxy ARP for a VLAN (i.e., IP subnetwork):

1. Click IP, ARP.
2. Select Configure General from the Step List.
3. Set the timeout to a suitable value for the ARP cache, or enable Proxy ARP for subnetworks that do not have routing or a default gateway.
4. Click Apply.

Figure 422: Configuring General Settings for ARP



The screenshot shows a web interface for configuring ARP settings. At the top, it says "Tools > ARP". Below that, there is a "Step:" dropdown menu currently set to "1. Configure General". Underneath, there is a "Timeout (300-86400)" label followed by a text input field containing the value "1200" and the unit "sec". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

Configuring Static ARP Addresses

For devices that do not respond to ARP requests or do not respond in a timely manner, traffic will be dropped because the IP address cannot be mapped to a physical address. If this occurs, use the IP > ARP (Configure Static Address – Add) page to manually map an IP address to the corresponding physical address in the ARP cache.

Command Usage

- ◆ The ARP cache is used to map 32-bit IP addresses into 48-bit hardware (that is, Media Access Control) addresses. This cache includes entries for hosts and other routers on local network interfaces defined on this router.
- ◆ You can define up to 256 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)

Web Interface

To map an IP address to the corresponding physical address in the ARP cache:

1. Click IP, ARP.
2. Select Configure Static Address from the Step List.
3. Select Add from the Action List.
4. Enter the IP address and the corresponding MAC address.
5. Click Apply.


Figure 423: Configuring Static ARP Entries

The screenshot shows a web interface for configuring static ARP entries. At the top, it says "Tools > ARP". Below this, there are two dropdown menus: "Step:" set to "2. Configure Static Address" and "Action:" set to "Add". The main area contains two input fields: "IP Address" with the value "10.2.78.105" and "MAC Address" with the value "00-e0-0c-00-df". To the right of the MAC Address field is a hint: "(xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx)". At the bottom right, there are two buttons: "Apply" and "Revert".

To display static entries in the ARP cache:

1. Click IP, ARP.
2. Select Configure Static Address from the Step List.
3. Select Show from the Action List.

Figure 424: Displaying Static ARP Entries



Displaying Dynamic or Local ARP Entries


Use the IP > ARP (Show Information – ARP Address) page to display dynamic or local entries in the ARP cache. The ARP cache contains static entries, and entries for local interfaces, including subnet, host, and broadcast addresses. However, most entries will be dynamically learned through replies to broadcast messages.

Web Interface

To display all dynamic and local entries in the ARP cache:

1. Click IP, ARP.
2. Select Show Information from the Step List.
3. Click ARP Address.

Figure 425: Displaying ARP Entries



Displaying ARP Statistics

Use the IP > ARP (Show Information – Statistics) page to display statistics for ARP messages crossing all interfaces on this router.

Parameters

These parameters are displayed:

Table 43: ARP Statistics

Parameter	Description
Received Request	Number of ARP Request packets received by the router.
Received Reply	Number of ARP Reply packets received by the router.

Table 43: ARP Statistics (Continued)

Parameter	Description
Sent Request	Number of ARP Request packets sent by the router.
Sent Reply	Number of ARP Reply packets sent by the router.

Web Interface

To display ARP statistics:

1. Click IP, ARP.
2. Select Show Information from the Step List.
3. Click Statistics.

Figure 426: Displaying ARP Statistics

Tool > ARP

Step: 2. Show Information

☐ ARP Address
☒ Statistics

ARP Statistics

Received Request	1234	Sent Request	43
Received Reply	2345	Sent Reply	432432

Refresh

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, or for creating an interface to multiple subnets. This switch supports both IPv4 and IPv6, and can be managed through either of these address types. For information on configuring the switch with an IPv6 address, see [“Setting the Switch's IP Address \(IP Version 6\)” on page 643](#).

Configuring IPv4 Interface Settings

Use the System > IP (Routing Interface – Add Address) page to configure an IPv4 address for the switch. An IPv4 address is obtained via DHCP by default for VLAN 1. To configure a static address, you need to change the switch's default settings to values that are compatible with your network. You may also need to establish a default gateway between the switch and management stations that exist on another network segment (if no routing protocols are enabled).

You can direct the device to obtain an address from a BOOTP or DHCP server, or manually configure a static IP address. Valid IP addresses consist of four decimal numbers, 0 to 255, separated by periods. Anything other than this format will not be accepted.

Command Usage

- ◆ This section describes how to configure a single local interface for initial access to the switch. To configure multiple IP interfaces, set up an IP interface for each VLAN.

- ◆ Once an IP address has been assigned to an interface, routing between different interfaces on the switch is enabled.
- ◆ To enable routing between interfaces defined on this switch and external network interfaces, you must configure static routes ([page 665](#)) or a default gateway ([page 664](#)).
- ◆ The precedence for configuring IP interfaces is the IP > General > Routing Interface (Add Address) menu, and then static routes ([page 665](#)).

Parameters

These parameters are displayed:

- ◆ **VLAN** – ID of the configured VLAN (1-4094). By default, all ports on the switch are members of VLAN 1. However, the management station can be attached to a port belonging to any VLAN, as long as that VLAN has been assigned an IP address.
- ◆ **IP Address Mode** – Specifies whether IP functionality is enabled via manual configuration (User Specified), Dynamic Host Configuration Protocol (DHCP), or Boot Protocol (BOOTP). If DHCP/BOOTP is enabled, IP will not function until a reply has been received from the server. Requests will be broadcast periodically by the switch for an IP address. DHCP/BOOTP responses can include the IP address, subnet mask, and default gateway. (Default: DHCP)
- ◆ **IP Address Type** – Specifies a primary or secondary IP address. An interface can have only one primary IP address, but can have many secondary IP addresses. In other words, secondary addresses need to be specified if more than one IP subnet can be accessed through this interface. For initial configuration, set this parameter to Primary. (Options: Primary, Secondary; Default: Primary)

Note that a secondary address cannot be configured prior to setting the primary IP address, and the primary address cannot be removed if a secondary address is still present. Also, if any router or switch in a network segment uses a secondary address, all other routers/switches in that segment must also use a secondary address from the same network or subnet address space.
- ◆ **IP Address** – IP Address of the VLAN. Valid IP addresses consist of four numbers, 0 to 255, separated by periods. (Default: None)



Note: You can manage the switch through any configured IP interface.

- ◆ **Subnet Mask** – This mask identifies the host address bits used for routing to specific subnets. (Default: None)
- ◆ **Restart DHCP** – Requests a new IP address from the DHCP server.

Web Interface

To set a static IPv4 address for the switch:

1. Click IP, General, Routing Interface.
2. Select Add Address from the Action list.
3. Select any configured VLAN, set IP Address Mode to "User Specified," set IP Address Type to "Primary" if no address has yet been configured for this interface, and then enter the IP address and subnet mask.
4. Select Primary or Secondary Address Type.
5. Click Apply.

Figure 427: Configuring a Static IPv4 Address

The screenshot shows the 'IP > General > Routing Interface' configuration page. The 'Action' dropdown is set to 'Add Address'. The 'VLAN' dropdown is set to '1'. The 'IP Address Mode' dropdown is set to 'User Specified'. The 'IP Address Type' dropdown is set to 'Primary'. The 'IP Address' text box contains '192.168.0.2' and the 'Subnet Mask' text box contains '255.255.255.0'. At the bottom, there is a 'Restart DHCP' button, a link 'Click this button to restart DHCP service.', and 'Apply' and 'Revert' buttons.

To obtain an dynamic IPv4 address through DHCP/BOOTP for the switch:

1. Click IP, General, Routing Interface.
2. Select Add Address from the Action list.
3. Select any configured VLAN, and set IP Address Mode to "BOOTP" or "DHCP."
4. Click Apply to save your changes.
5. Then click Restart DHCP to immediately request a new address.

IP will be enabled but will not function until a BOOTP or DHCP reply is received. Requests are broadcast every few minutes using exponential backoff until IP configuration information is obtained from a BOOTP or DHCP server.

Figure 428: Configuring a Dynamic IPv4 Address

IP > General > Routing Interface

Action: Add Address

VLAN: 1

IP Address Mode: DHCP

IP Address Type: Primary

IP Address:

Subnet Mask:

Restart DHCP [Click this button to restart DHCP service.](#)

Apply Revert



Note: The switch will also broadcast a request for IP configuration settings on each power reset.

Note: If you lose the management connection, make a console connection to the switch and enter “show ip interface” to determine the new switch address.

Renewing DHCP – DHCP may lease addresses to clients indefinitely or for a specific period of time. If the address expires or the switch is moved to another network segment, you will lose management access to the switch. In this case, you can reboot the switch or submit a client request to restart DHCP service via the CLI.

If the address assigned by DHCP is no longer functioning, you will not be able to renew the IP settings via the web interface. You can only restart DHCP service via the web interface if the current address is still available.

To show the IPv4 address configured for an interface:

1. Click IP, General, Routing Interface.
2. Select Show Address from the Action list.
3. Select an entry from the VLAN list.

Figure 429: 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	192.168.0.2	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 639](#).

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 430: Configuring the IPv6 Default Gateway



The screenshot shows a web interface for configuring IPv6 settings. At the top, it says "IP > IPv6". Below that, there is a dropdown menu labeled "Action:" with "Configure Global" selected. Underneath, there is a field labeled "Default Gateway" with the value "2001:DB8:2222:7272::254" entered. At the bottom right, there are two buttons: "Apply" and "Revert".

Configuring IPv6 Interface Settings

Use the IP > IPv6 Configuration (Configure Interface) pages 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 649](#)).

- ◆ IPv6 Neighbor Discovery Protocol supersedes IPv4 Address Resolution Protocol in IPv6 networks. IPv6 nodes on the same network segment use Neighbor Discovery to discover each other's presence, to determine each other's link-layer addresses, to find routers and to maintain reachability information about the paths to active neighbors. The key parameters used to facilitate this process are the number of attempts made to verify whether or not a duplicate address exists on the same network segment, and the interval between neighbor solicitations used to verify reachability information.

Parameters

These parameters are displayed:

VLAN Mode

- ◆ **VLAN** – ID of a configured VLAN which is to be used for management access, or as a standard interface for a subnet. By default, all ports on the switch are members of VLAN 1. However, the management station can be attached to a port belonging to any VLAN, as long as that VLAN has been assigned an IP address. (Range: 1-4094)
- ◆ **Enable IPv6 Explicitly** – Enables IPv6 on an interface and assigns it a link-local address. Note that when an explicit address is assigned to an interface, IPv6 is automatically enabled, and cannot be disabled until all assigned addresses have been removed. (Default: Disabled)

Disabling this parameter does not disable IPv6 for an interface that has been explicitly configured with an IPv6 address.
- ◆ **MTU** – Sets the size of the maximum transmission unit (MTU) for IPv6 packets sent on an interface. (Range: 1280-65535 bytes; Default: 1500 bytes)
 - The maximum value set in this field cannot exceed the MTU of the physical interface, which is currently fixed at 1500 bytes.
 - If a non-default value is configured, an MTU option is included in the router advertisements sent from this device. This option is provided to ensure that all nodes on a link use the same MTU value in cases where the link MTU is not otherwise well known.
 - IPv6 routers do not fragment IPv6 packets forwarded from other routers. However, traffic originating from an end-station connected to an IPv6 router may be fragmented.
 - All devices on the same physical medium must use the same MTU in order to operate correctly.
 - IPv6 must be enabled on an interface before the MTU can be set. If an IPv6 address has not been assigned to the switch, "N/A" is displayed in the MTU field.
- ◆ **ND DAD Attempts** – The number of consecutive neighbor solicitation messages sent on an interface during duplicate address detection. (Range: 0-600, Default: 3)

- 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 163](#)). 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 switch.
- ◆ **Restart DHCPv6** – When DHCPv6 is restarted, the switch may attempt to acquire an IP address prefix through stateful address autoconfiguration. If the router advertisements have the “other stateful configuration” flag set, the switch may also attempt to acquire other non-address configuration information (such as a default gateway) when DHCPv6 is restarted.

Prior to submitting a client request to a DHCPv6 server, the switch should be configured with a link-local address using the Address Autoconfig option. The state of the Managed Address Configuration flag (M flag) and Other Stateful Configuration flag (O flag) received in Router Advertisement messages will determine the information this switch should attempt to acquire from the DHCPv6 server as described below.

- Both M and O flags are set to 1:
DHCPv6 is used for both address and other configuration settings.
This combination is known as DHCPv6 stateful autoconfiguration, in which a DHCPv6 server assigns stateful addresses to IPv6 hosts.
- The M flag is set to 0, and the O flag is set to 1:
DHCPv6 is used only for other configuration settings.
Neighboring routers are configured to advertise non-link-local address prefixes from which IPv6 hosts derive stateless addresses.
This combination is known as DHCPv6 stateless autoconfiguration, in which a DHCPv6 server does not assign stateful addresses to IPv6 hosts, but does assign stateless configuration settings.

RA Guard

- ◆ **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 431: Configuring General Settings for an IPv6 Interface

The screenshot shows the 'IP > IPv6 Configuration' web interface. At the top, the 'Action:' dropdown is set to 'Configure Interface'. Below this, the 'Mode' section has two radio buttons: 'VLAN' (which is selected) and 'RA Guard'. Under the 'VLAN' mode, there is a 'VLAN' dropdown menu currently showing '1'. Below the VLAN dropdown, there is a checkbox for 'Enable IPv6 Explicitly' which is currently unchecked. Further down, there are five input fields: 'MTU (1280-65535)' with a value of '1500' and 'bytes' unit; 'ND DAD Attempts (0-600)' with a value of '1'; 'ND NS Interval (1000-3600000)' with a value of '1000' and 'ms' unit; and 'ND Reachable-Time (0-3600000)' with a value of '30000' and 'ms' unit. At the bottom right of the form, there are two buttons: 'Apply' and 'Revert'.

To configure RA Guard for the switch:

1. Click IP, IPv6 Configuration.
2. Select Configure Interface from the Action list.
3. Select RA Guard mode.
4. Enable RA Guard for untrusted interfaces.
5. Click Apply.

Figure 432: Configuring RA Guard for an IPv6 Interface

The screenshot shows the 'IP > IPv6 Configuration' page. The 'Action' dropdown is set to 'Configure Interface'. Under 'Mode', 'RA Guard' is selected. Under 'Interface', 'Port' is selected. A 'Port List' table shows ports 1 through 5, with a total of 28 ports. The 'RA Guard' column for each port has a checkbox labeled 'Enabled'. Port 5's checkbox is checked.

Port	RA Guard
1	☐ Enabled
2	☐ Enabled
3	☐ Enabled
4	☐ Enabled
5	☒ Enabled

Configuring an IPv6 Address Use the IP > IPv6 Configuration (Add IPv6 Address) page to configure an IPv6 interface for management access over the network, or for creating an interface to multiple subnets.

Command Usage

- ◆ All IPv6 addresses must be formatted according to RFC 2373 “IPv6 Addressing Architecture,” using 8 colon-separated 16-bit hexadecimal values. One double colon may be used in the address to indicate the appropriate number of zeros required to fill the undefined fields.
- ◆ The switch must 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 644](#)), 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 644](#)).
 - It can be manually configured by specifying the entire network prefix and prefix length, and using the EUI-64 form of the interface identifier to automatically create the low-order 64 bits in the host portion of the address.

- You can also manually configure the global unicast address by entering the full address and prefix length.
- ◆ You can configure multiple IPv6 global unicast addresses per interface, but only one link-local address per interface.
- ◆ If a duplicate link-local address is detected on the local segment, this interface is disabled and a warning message displayed on the console. If a duplicate global unicast address is detected on the network, the address is disabled on this interface and a warning message displayed on the console.
- ◆ When an explicit address is assigned to an interface, IPv6 is automatically enabled, and cannot be disabled until all assigned addresses have been removed.

Parameters

These parameters are displayed:

- ◆ **VLAN** – ID of a configured VLAN which is to be used for management access, or for creating an interface to multiple subnets. By default, all ports on the switch are members of VLAN 1. However, the management station can be attached to a port belonging to any VLAN, as long as that VLAN has been assigned an IP address. (Range: 1-4094)
- ◆ **Address Type** – Defines the address type configured for this interface.
 - **Global** – Configures an IPv6 global unicast address with a full IPv6 address including the network prefix and host address bits, followed by a forward slash, and a decimal value indicating how many contiguous bits (from the left) of the address comprise the prefix (i.e., the network portion of the address).
 - **EUI-64** (Extended Universal Identifier) – Configures an IPv6 address for an interface using an EUI-64 interface ID in the low order 64 bits.
 - When using EUI-64 format for the low-order 64 bits in the host portion of the address, the value entered in the IPv6 Address field includes the network portion of the address, and the prefix length indicates how many contiguous bits (starting at the left) of the address comprise the prefix (i.e., the network portion of the address). Note that the value specified in the IPv6 Address field may include some of the high-order host bits if the specified prefix length is less than 64 bits. If the specified prefix length exceeds 64 bits, then the bits used in the network portion of the address will take precedence over the interface identifier.
 - IPv6 addresses are 16 bytes long, of which the bottom 8 bytes typically form a unique host identifier based on the device's MAC address. The EUI-64 specification is designed for devices that use an extended 8-byte MAC address. For devices that still use a 6-byte MAC address (also known as EUI-48 format), it must be converted into EUI-64 format by

inverting the universal/local bit in the address and inserting the hexadecimal number FFFE between the upper and lower three bytes of the MAC address.

For example, if a device had an EUI-48 address of 28-9F-18-1C-82-35, the global/local bit must first be inverted to meet EUI-64 requirements (i.e., 1 for globally defined addresses and 0 for locally defined addresses), changing 28 to 2A. Then the two bytes FFFE are inserted between the OUI (i.e., organizationally unique identifier, or company identifier) and the rest of the address, resulting in a modified EUI-64 interface identifier of 2A-9F-18-FF-FE-1C-82-35.

- This host addressing method allows the same interface identifier to be used on multiple IP interfaces of a single device, as long as those interfaces are attached to different subnets.
- **Link Local** – Configures an IPv6 link-local address.
 - The address prefix must be in the range of FE80~FEBF.
 - You can configure only one link-local address per interface.
 - The specified address replaces a link-local address that was automatically generated for the interface.

◆ **IPv6 Address** – IPv6 address assigned to this interface.

Web Interface

To configure an IPv6 address:

1. Click IP, IPv6 Configuration.
2. Select Add IPv6 Address from the Action list.
3. Specify the VLAN to configure, select the address type, and then enter an IPv6 address and prefix length.
4. Click Apply.

Figure 433: Configuring an IPv6 Address

The screenshot shows a web interface for configuring an IPv6 address. The breadcrumb is "IP > IPv6". The "Action:" dropdown is set to "Add IPv6 Address". Below this, there are three rows of configuration options: "VLAN" with a dropdown set to "1", "Address Type" with a dropdown set to "Global", and "IPv6 Address" with a text input field containing "2001:DB8:2222:7272::72/96". At the bottom right, there are two buttons: "Apply" and "Revert".

Showing IPv6 Addresses Use the IP > IPv6 Configuration (Show IPv6 Address) page to display the IPv6 addresses assigned to an interface.

Parameters

These parameters are displayed:

- ◆ **VLAN** – ID of a configured VLAN. By default, all ports on the switch are members of VLAN 1. However, the management station can be attached to a port belonging to any VLAN, as long as that VLAN has been assigned an IP address. (Range: 1-4094)
- ◆ **IPv6 Address Type** – The address type (Global, EUI-64, Link Local).
- ◆ **IPv6 Address** – An IPv6 address assigned to this interface.

In addition to the unicast addresses assigned to an interface, a node is also required to listen to the all-nodes multicast addresses FF01::1 (interface-local scope) and FF02::1 (link-local scope).

FF01::1/16 is the transient interface-local multicast address for all attached IPv6 nodes, and FF02::1/16 is the link-local multicast address for all attached IPv6 nodes. The interface-local multicast address is only used for loopback transmission of multicast traffic. Link-local multicast addresses cover the same types as used by link-local unicast addresses, including all nodes (FF02::1), all routers (FF02::2), and solicited nodes (FF02::1:FFXX:XXXX) as described below.

A node is also required to compute and join the associated solicited-node multicast addresses for every unicast and anycast address it is assigned. IPv6 addresses that differ only in the high-order bits, e.g. due to multiple high-order prefixes associated with different aggregations, will map to the same solicited-node address, thereby reducing the number of multicast addresses a node must join. In this example, FF02::1:FF90:0/104 is the solicited-node multicast address which is formed by taking the low-order 24 bits of the address and appending those bits to the prefix.

Note that the solicited-node multicast address (link-local scope FF02) is used to resolve the MAC addresses for neighbor nodes since IPv6 does not support the broadcast method used by the Address Resolution Protocol in IPv4.

These additional addresses are displayed by the “show ip interface” command described in the *CLI Reference Guide*.

- ◆ **Configuration Mode** – Indicates if this address was automatically generated or manually configured.


Web Interface

To show the configured IPv6 addresses:

1. Click IP, IPv6 Configuration.
2. Select Show IPv6 Address from the Action list.

3. Select a VLAN from the list.

Figure 434: 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 44: Show IPv6 Neighbors - display description

Field	Description
IPv6 Address	IPv6 address of neighbor.
Age	The time since the address was verified as reachable (in seconds). A static entry is indicated by the value "Permanent."
Link-layer Address	Physical layer MAC address.
State	The following states are used for dynamic entries: ◆ 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.

Table 44: Show IPv6 Neighbors - display description (Continued)


Field	Description
	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 435: Showing IPv6 Neighbors



The screenshot shows the 'IP > IPv6 Configuration' page. Under the 'Action:' dropdown, 'Show IPv6 Neighbor Cache' is selected. Below this, a section titled 'Current Neighbor Cache Table' shows 'Total: 1'. A table displays the neighbor cache information:

IPv6 Address	Age	Link-layer Address	State	VLAN
2001:DB8:2222:7272::74	20	00-E0-0C-9C-CA-10	Reachable	1

A 'Clear' button is located at the bottom right of the table.

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 45: Show IPv6 Statistics - display description

Field	Description
IPv6 Statistics	
<i>IPv6 Received</i>	
Total	The total number of input datagrams received by the interface, including those received in error.
Header Errors	The number of input datagrams discarded due to errors in their IPv6 headers, including version number mismatch, other format errors, hop count exceeded, IPv6 options, etc.
Too Big Errors	The number of input datagrams that could not be forwarded because their size exceeded the link MTU of outgoing interface.
No Routes	The number of input datagrams discarded because no route could be found to transmit them to their destination.
Address Errors	The number of input datagrams discarded because the IPv6 address in their IPv6 header's destination field was not a valid address to be received at this entity. This count includes invalid addresses (e.g., ::0) and unsupported addresses (e.g., addresses with unallocated prefixes). For entities which are not IPv6 routers and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
Unknown Protocols	The number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol. This counter is incremented at the interface to which these datagrams were addressed which might not be necessarily the input interface for some of the datagrams.
Truncated Packets	The number of input datagrams discarded because datagram frame didn't carry enough data.
Discards	The number of input IPv6 datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (e.g., for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting re-assembly.
Delivers	The total number of datagrams successfully delivered to IPv6 user-protocols (including ICMP). This counter is incremented at the interface to which these datagrams were addressed which might not be necessarily the input interface for some of the datagrams.
Reassembly Request Datagrams	The number of IPv6 fragments received which needed to be reassembled at this interface. Note that this counter is incremented at the interface to which these fragments were addressed which might not be necessarily the input interface for some of the fragments.

Table 45: Show IPv6 Statistics - display description (Continued)

Field	Description
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 ipv6IfStatsOutForwDatagrams.
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 ipv6IfStatsOutForwDatagrams if any such packets met this (discretionary) discard criterion.
No Routes	The number of input datagrams discarded because no route could be found to transmit them to their destination.
Generated Fragments	The number of output datagram fragments that have been generated as a result of fragmentation at this output interface.
Fragment Succeeded	The number of IPv6 datagrams that have been successfully fragmented at this output interface.
Fragment Failed	The number of IPv6 datagrams that have been discarded because they needed to be fragmented at this output interface but could not be.
ICMPv6 Statistics	
<i>ICMPv6 received</i>	
Input	The total number of ICMP messages received by the interface which includes all those counted by ipv6IfIcmpInErrors. 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.

Table 45: Show IPv6 Statistics - display description (Continued)

Field	Description
Parameter Problem Messages	The number of ICMP Parameter Problem messages received by the interface.
Echo Request Messages	The number of ICMP Echo (request) messages received by the interface.
Echo Reply Messages	The number of ICMP Echo Reply messages received by the interface.
Router Solicit Messages	The number of ICMP Router Solicit messages received by the interface.
Router Advertisement Messages	The number of ICMP Router Advertisement messages received by the interface.
Neighbor Solicit Messages	The number of ICMP Neighbor Solicit messages received by the interface.
Neighbor Advertisement Messages	The number of ICMP Neighbor Advertisement messages received by the interface.
Redirect Messages	The number of Redirect messages received by the interface.
Group Membership Query Messages	The number of ICMPv6 Group Membership Query messages received by the interface.
Group Membership Response Messages	The number of ICMPv6 Group Membership Response messages received by the interface.
Group Membership Reduction Messages	The number of ICMPv6 Group Membership Reduction messages received by the interface.
Multicast Listener Discovery Version 2 Reports	The number of MLDv2 reports received by the interface.
<i>ICMPv6 Transmitted</i>	
Output	The total number of ICMP messages which this interface attempted to send. Note that this counter includes all those counted by icmpOutErrors.
Destination Unreachable Messages	The number of ICMP Destination Unreachable messages sent by the interface.
Packet Too Big Messages	The number of ICMP Packet Too Big messages sent by the interface.
Time Exceeded Messages	The number of ICMP Time Exceeded messages sent by the interface.
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.

Table 45: Show IPv6 Statistics - display description (Continued)

Field	Description
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 436: Showing IPv6 Statistics (IPv6)

The screenshot shows a web interface for configuring IPv6. At the top, there's a breadcrumb 'IP > IPv6'. Below it, an 'Action:' dropdown menu is set to 'Show Statistics'. Underneath, there are three radio buttons for 'Type': 'IPv6' (selected), 'ICMPv6', and 'UDP'. The main section is titled 'IPv6 Statistics' and contains a table of statistics. The table has two columns: the statistic name and its value. The statistics are: Total Received (55), Received Header Errors (0), Received Too Big Errors (0), Received No Routes (0), Received Address Errors (0), Received Unknown Protocols (0), Received Truncated Packets (0), Received Discards (0), Received Delivers (55), Received Reassembly Request Datagrams (0), Received Reassembled Succeeded, Received Reassembled Failed, Transmitted Forwards Datagrams, Transmitted Requests, Transmitted Discards, Transmitted No Routes, Transmitted Generated Fragments, Transmitted Fragment Succeeded, and Transmitted Fragment Failed. A 'Clear' button is located at the bottom right of the statistics table.

IPv6 Statistics	
Total Received	55
Received Header Errors	0
Received Too Big Errors	0
Received No Routes	0
Received Address Errors	0
Received Unknown Protocols	0
Received Truncated Packets	0
Received Discards	0
Received Delivers	55
Received Reassembly Request Datagrams	0
Received Reassembled Succeeded	
Received Reassembled Failed	
Transmitted Forwards Datagrams	
Transmitted Requests	
Transmitted Discards	
Transmitted No Routes	
Transmitted Generated Fragments	
Transmitted Fragment Succeeded	
Transmitted Fragment Failed	

Figure 437: Showing IPv6 Statistics (ICMPv6)

IP > IPv6 Configuration

Action: Show Statistics

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 Message	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 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		

Clear

Figure 438: Showing IPv6 Statistics (UDP)

IP > IPv6

Action: Show Statistics

Type: ☐ IPv6 ☐ ICMPv6 ☒ UDP

UDP Statistics

Input	10
No Port Errors	0
Other Errors	0
Output	1

Clear

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 46: 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 439: Showing Reported MTU Values



The screenshot shows the 'IP > IPv6' configuration page. Under the 'Action:' dropdown, 'Show MTU' is selected. Below this, a table titled 'MTU Table Total: 2' displays the following data:

MTU	Since	Destination Address
1400	00:04:21	5000:1::3
1280	00:04:50	FE80::203:A0FF:FED6:141D

This chapter provides information on network functions including:

- ◆ [Static Routes](#) – Configures static routes to other network segments.
- ◆ [Routing Table](#) – Displays routing entries learned through dynamic routing and 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 via the default gateway or static routes. 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 163](#)), assign all ports that belong to the same group to these VLANs ([page 165](#)), and then assign an IP interface to each VLAN ([page 639](#) or [page 643](#)). 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 440: 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 IP Routing Interfaces

Configuring Local and Remote Interfaces Use the IP > General > Routing Interface (Add Address) page to configure routing interfaces for directly connected IPv4 subnets (see [“Setting the Switch’s IP Address \(IP Version 4\)” on page 639](#)). Or use the IP > IPv6 Configuration pages to configure routing interfaces for directly connected IPv6 subnets (see [“Setting the Switch’s IP Address \(IP Version 6\)” on page 643](#)).

If this router is directly connected to end node devices (or connected to end nodes through shared media) that will be assigned to a specific subnet, then you must create a router interface for each VLAN that will support routing. The router interface consists of an IP address and subnet mask. This interface address defines both the network prefix number to which the router interface is attached and the router’s host number on that network. In other words, a router interface address defines the network segment that is connected to that interface, and allows you to send IP packets to or from the router.

You can specify the IP subnets connected directly to this router by manually assigning an IP address to each VLAN, or using BOOTP or DHCP to dynamically assign an address. To specify IP subnets not directly connected to this router, you can configure static routes (see [page 665](#)).

Once IP interfaces have been configured, the switch functions as a multilayer routing switch, operating at either Layer 2 or 3 as required. All IP packets are routed directly between local interfaces, or indirectly to remote interfaces using static routing. All other packets for non-IP protocols (for example, NetBuei, NetWare or AppleTalk) are switched based on MAC addresses).

To route traffic between remote IP interfaces, the switch should be recognized by other network nodes as an IP router, either by setting it to advertise itself as the default gateway or by redirection from another router via the ICMP process used by various routing protocols.

If another router is designated as the default gateway, then this switch will pass packets to the router for any unknown hosts or subnets.

To configure the default gateway for IPv4, use the static routing table as described on [page 665](#), enter 0.0.0.0 for the IP address and subnet mask, and then specify this switch itself or another router as the gateway. To configure a gateway for IPv6, see [“Configuring the IPv6 Default Gateway” on page 643](#).

Configuring Static Routes

You can enter static routes in the routing table using the IP > Routing > Static Routes (Add) page. Static routes can be set 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 256 static routes can be configured.

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 used for route selection. If more than one route is configured for the same destination, the route with the longest matching prefix is used. If more than one route is configured to the same destination with the same matching prefix, then the route with the smallest distance is used. (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 441: Configuring Static Routes

IP > Routing > Static Routes

Action: Add ▾

Destination IP Address

10.2.48.0

Net Mask / Prefix Length

255.255.255.0

Next Hop

10.2.48.1

Distance (1-255)

5

(Optional)

Apply

Revert

To display static routes:

1. Click IP, Routing, Static Routes.
2. Select Show from the Action List.

Figure 442: Displaying Static Routes

IP > Routing > Static Routes

Action: Show ▾

Static Table List Total: 1

	Destination IP Address	Net Mask / Prefix Length	Next Hop	Distance
☐	10.2.48.0	255.255.255.0	10.2.48.1	5

Delete

Revert

Displaying the Routing Table

Use the IP > Routing > Routing Table (Show Information) page to display all routes that can be accessed via local network interfaces or 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 static route entry must be up. Note that routes currently not accessible for forwarding, may still be displayed by using the “show ip route database” command described in the *CLI Reference Guide*.

Parameters

These parameters are displayed:

- ◆ **VLAN** – VLAN identifier (i.e., configured as a valid IP subnet).
- ◆ **Destination IP Address** – IP address of the destination network, subnetwork, or host. Note that the address 0.0.0.0 indicates the default gateway for this router.
- ◆ **Net Mask / Prefix Length** – Network mask for the associated IP subnet. This mask identifies the host address bits used for routing to specific subnets.
- ◆ **Next Hop** – The IP address of the next hop (or gateway) in this route.
- ◆ **Metric** – Cost for this interface.
- ◆ **Protocol** – The protocol which generated this route information. (Options: Local, Static, Others)

Web Interface

To display the routing table:

1. Click IP, Routing, Routing Table.

Figure 443: Displaying the Routing Table

The screenshot shows a web interface for displaying the routing table. At the top, it says "IP > Routing > Routing Table". Below this is an "Action:" dropdown menu set to "Show Information". Underneath is a section titled "Routing Table List" with a "Total: 5" count. The main part of the interface is a table with the following data:

VLAN	Destination IP Address	Net Mask / Prefix Length	Next Hop	Metric	Protocol
0	127.0.0.0	255.0.0.0	--	0	Local
2	192.168.0.0	255.255.255.0	--	0	Local
1	192.168.2.0	255.255.255.0	--	0	Local
2	192.168.3.0	255.255.255.0	192.168.0.1	0	Static
0	:::1	128	--	0	Local

This chapter describes the following IP services:

- ◆ **DNS** – Configures default domain names, identifies servers to use for dynamic lookup, and shows how to configure static entries.
- ◆ **DHCP Client** – Specifies the DHCP client identifier for an interface.
- ◆ **DHCP Relay** – Enables DHCP relay service, and defines the servers to which client requests are forwarded.
- ◆ **DHCP Dynamic Provision** – Enables dynamic provision via DHCP.
- ◆ **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 672](#)).
- ◆ 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 444: Configuring General Settings for DNS

IP Service > DNS > General

Action: Configure Global

Domain Lookup ☒ Enabled

Default Domain Name

Apply 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 669](#)). 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 672](#)).
- ◆ If all name servers are deleted, DNS will automatically be disabled.

Parameters

These parameters are displayed:

Domain Name – Name of the host. Do not include the initial dot that separates the host name from the domain name. (Range: 1-68 characters)

Web Interface

To create a list domain names:

1. Click IP Service, DNS.
2. Select Add Domain Name from the Action list.
3. Enter one domain name at a time.
4. Click Apply.

Figure 445: Configuring a List of Domain Names for DNS

The screenshot shows the 'IP Service > DNS > General' configuration page. The 'Action:' dropdown menu is set to 'Add Domain Name'. Below it, the 'Domain Name' text field contains 'sample.com.uk'. At the bottom right, there are 'Apply' and 'Revert' buttons.

To show the list domain names:

1. Click IP Service, DNS.
2. Select Show Domain Names from the Action list.

Figure 446: Showing the List of Domain Names for DNS

The screenshot shows the 'IP Service > DNS > General' configuration page. The 'Action:' dropdown menu is set to 'Show Domain Names'. Below it, the 'Domain Name List' section shows a table with two entries: 'google.com' and 'hinet.net'. The 'Total: 2' is displayed next to the list. At the bottom right, there are 'Delete' and 'Revert' buttons.

Domain Name
google.com
hinet.net

Configuring a List of Name Servers Use the IP Service > DNS - General (Add Name Server) page to configure a list of name servers to be tried in sequential order.

Command Usage

- ◆ To enable DNS service on this switch, configure one or more name servers, and enable domain lookup status (see [“Configuring General DNS Service Parameters” on page 669](#)).
- ◆ 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 447: 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 448: 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 449: Configuring Static Entries in the DNS Table

IP Service > DNS > Static Host Table

Action: Add

Host Name

IP Address

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 450: 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	124.29.31.155
☐	google.com	133.45.211.18

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 451: Showing Entries in the DNS Cache



The screenshot shows a web interface titled "IP Service > DNS > Cache". Below the title bar, it says "Cache Information Total: 3". There is a table with 6 columns: No., Flag, Type, IP, TTL, and Host. The table contains 3 rows of data. Below the table is a "Clear" button.

No.	Flag	Type	IP	TTL	Host
1	4	CNAME	192.168.110.2	360	www.sina.com.cn
2	4	CNAME	10.2.44.3	892	www.yahoo.akadns.new
3	4	ALIAS	pointer to: 2	298	www.yahoo.com

Dynamic Host Configuration Protocol

Dynamic Host Configuration Protocol (DHCP) can dynamically allocate an IP address and other configuration information to network clients when they boot up. If a subnet does not already include a BOOTP or DHCP server, you can relay DHCP client requests to a DHCP server on another subnet.

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 47: 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 48: Options 55 and 124 Statements

Option	Statement	
	Keyword	Parameter
55	dhcp-parameter-request-list	a list of parameters, separated by “
124	vendor-class-identifier	a string indicating the vendor class identifier

- ◆ The server should reply with the TFTP server name and boot file name.
- ◆ Note that the vendor class identifier can be formatted in either text or hexadecimal, but the format used by both the client and server must be the same.

Parameters

These parameters are displayed:

- ◆ **VLAN** – ID of configured VLAN.
- ◆ **Vendor Class ID** – The following options are supported when the check box is marked to enable this feature:
 - **Default** – The default string is the model number.
 - **Text** – A text string. (Range: 1-32 characters)
 - **Hex** – A hexadecimal value.

Web Interface

To configure a DHCP client identifier:

1. Click IP Service, DHCP, Client.
2. Mark the check box to enable this feature. Select the default setting, or the format for a vendor class identifier. If a non-default value is used, enter a text string or hexadecimal value.
3. Click Apply.

Figure 452: Specifying A DHCP Client Identifier

IP Service > DHCP > Client

VLAN: 1


Vendor Class ID: ☒ Default

Apply Revert

Configuring DHCP Relay Service

Use the IP Service > DHCP > Relay page to configure DHCP relay service for attached host devices. If DHCP relay is enabled, and this switch sees a DHCP request broadcast, it inserts its own IP address into the request so that the DHCP server will know the subnet where the client is located. Then, the switch forwards the packet to the DHCP server. When the server receives the DHCP request, it allocates a free IP address for the DHCP client from its defined scope for the DHCP client's subnet, and sends a DHCP response back to the DHCP relay agent (i.e., this switch). This switch then passes the DHCP response received from the server to the client.

Figure 453: Layer 3 DHCP Relay Service



Command Usage

- ◆ You must specify the IP address for at least one active DHCP server. Otherwise, the switch's DHCP relay agent will not be able to forward client requests to a DHCP server. Up to five DHCP servers can be specified in order of preference.
- ◆ DHCP relay configuration will be disabled if an active DHCP server is detected on the same network segment.

Parameters

These parameters are displayed:

- ◆ **VLAN ID** – ID of configured VLAN.
- ◆ **Server IP Address** – Addresses of DHCP servers or relay servers to be used by the switch's DHCP relay agent in order of preference.
- ◆ **Restart DHCP Relay** – Use this button to re-initialize DHCP relay service.

Web Interface

To configure DHCP relay service:

1. Click IP Service, DHCP, Relay.
2. Enter up to five IP addresses for DHCP servers or relay servers in order of preference for any VLAN.
3. Click Apply.

Figure 454: Configuring DHCP Relay Service

IP Service > DHCP > Relay

Note: DHCP relay configuration will be disabled if an active DHCP server is detected on the same network segment.

DHCP Server by VLAN List Total: 1

VLAN	Server IP Address
1	192.168.2.33

Restart DHCP Relay Click the button to restart DHCP Relay service.

Apply Revert

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 bootstrap.
3. Click Apply.

Figure 455: Enabling Dynamic Provisioning via DHCP



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 456: Configuring Global Settings for PPPoE Intermediate Agent

IP Service > PPPoE Intermediate Agent

Step: 1. Configure Global

PPPoE IA Global Status	☐ Enabled
Access Node Identifier	
Operational Access Node Identifier	10.254.254.1
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.
- ◆ **Operational Circuit ID** – The configured circuit identifier.

- ◆ **Remote ID** – String identifying the remote identifier (or interface) on this switch to which the user is connected. (Range: 1-63 ASCII characters; Default: Port MAC address)
- ◆ **Operational Remote ID** – The configured circuit identifier.

Web Interface

To configure interface settings for PPPoE IA:

1. Click IP Service, PPPoE Intermediate Agent.
2. Select Configure Interface from the Step list.
3. Select Port or Trunk interface type.
4. Enable PPPoE IA on an interface, set trust status, enable vendor tag stripping if required, and set the circuit ID and remote ID.
5. Click Apply.

Figure 457: Configuring Interface Settings for PPPoE Intermediate Agent

Port	PPPoE IA Status	Trust Status	Vendor Tag Strip	Circuit ID	Operation Circuit ID	Remote ID	Operation Remote ID
1	☐ Enabled	☐ Enabled	☐ Enabled		1/1:vid		00-E0-0C-00-00-FE
2	☐ Enabled	☐ Enabled	☐ Enabled		1/2:vid		00-E0-0C-00-00-FF
3	☐ Enabled	☐ Enabled	☐ Enabled		1/3:vid		00-E0-0C-00-01-00
4	☐ Enabled	☐ Enabled	☐ Enabled		1/4:vid		00-E0-0C-00-01-01
5	☐ Enabled	☐ Enabled	☐ Enabled		1/5:vid		00-E0-0C-00-01-02

Showing PPPoE IA Statistics Use the IP Service > PPPoE Intermediate Agent (Show Statistics) page to show statistics on PPPoE IA protocol messages.

Parameters

These parameters are displayed:

- ◆ **Interface** – Port or trunk selection.
- ◆ **Received** – Received PPPoE active discovery messages.
 - **All** – All PPPoE active discovery message types.
 - **PADI** – PPPoE Active Discovery Initiation messages.
 - **PADO** – PPPoE Active Discovery Offer messages.

- **PADR** – PPPoE Active Discovery Request messages.
- **PADS** – PPPoE Active Discovery Session-Confirmation messages.
- **PADT** – PPPoE Active Discovery Terminate messages.
- ◆ **Dropped** – Dropped PPPoE active discovery messages.
 - **Response from untrusted** – Response from an interface which not been configured as trusted.
 - **Request towards untrusted** – Request sent to an interface which not been configured as trusted.
 - **Malformed** – Corrupted PPPoE message.

Web Interface

To show statistics for PPPoE IA protocol messages:

1. Click IP Service, PPPoE Intermediate Agent.
2. Select Show Statistics from the Step list.
3. Select Port or Trunk interface type.

Figure 458: Showing PPPoE Intermediate Agent Statistics

IP Service > PPPoE Intermediate Agent

Step: 3. Show Statistics

Interface ☒ Port ☐ Trunk

PPPoE Intermediate Agent Statistics Total: 12

	Port	Received						Dropped		
		All	PADI	PADO	PADR	PADS	PADT	Response from untrusted	Request towards untrusted	Malformed
☐	1	0	0	0	0	0	0	0	0	0
☐	2	0	0	0	0	0	0	0	0	0
☐	3	0	0	0	0	0	0	0	0	0
☐	4	0	0	0	0	0	0	0	0	0
☐	5	0	0	0	0	0	0	0	0	0

Section III

Appendices

This section provides additional information and includes these items:

- ◆ [“Software Specifications” on page 687](#)
- ◆ [“Troubleshooting” on page 693](#)
- ◆ [“License Information” on page 695](#)



Software Specifications

Software Features

Management Authentication Local, RADIUS, TACACS+, Port Authentication (802.1X), HTTPS, SSH, Port Security, IP Filter

Client Access Control Access Control Lists (2048 rules), Port Authentication (802.1X), MAC Authentication, Port Security, DHCP Snooping, IP Source Guard

Port Configuration 1000BASE-T: 10/100 Mbps at half/full duplex, 1000 Mbps at full duplex
1000BASE-SX/LX - 1000 Mbps full duplex (SFP)
10GBASE-CR/SR/LR/ER - 10 Gbps full duplex (SFP+)

Flow Control Full Duplex: IEEE 802.3-2005
Half Duplex: Back pressure

Storm Control Broadcast, multicast, or unicast traffic throttled above a critical threshold

Port Mirroring 1 session, one or more source ports to one analysis 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), private VLANs, 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 combination of strict and weighted queueing
Layer 3/4 priority mapping: IP Port, IP Precedence, IP DSCP

Quality of Service DiffServ supports class maps, policy maps, and service policies

Multicast Filtering IGMP Snooping (Layer 2 IPv4)
MLD Snooping (Layer 2 IPv6)
IGMP (Layer 3)
Multicast VLAN Registration (IPv4/IPv6)

IP Routing ARP, Proxy ARP
Static routes
CIDR (Classless Inter-Domain Routing)

Additional Features BOOTP Client
Connectivity Fault Management
DHCP Client, Relay, Option 82
DNS Client, Proxy
ERPS (Ethernet Ring Protection Switching)
LLDP (Link Layer Discover Protocol)
OAM (Operation, Administration, and Maintenance)
RMON (Remote Monitoring, groups 1,2,3,9)
SMTP Email Alerts
SNMP (Simple Network Management Protocol)
SNTP (Simple Network Time Protocol)

Management Features

In-Band Management Telnet, web-based HTTP or HTTPS, SNMP manager, or Secure Shell

Out-of-Band Management RS-232 DB-9 console port

Software Loading HTTP, FTP or TFTP in-band, or XModem out-of-band

SNMP Management access via MIB database
Trap management to specified hosts

RMON Groups 1, 2, 3, 9 (Statistics, History, Alarm, Event)

Standards

Ethernet Service OAM (ITU-T Y.1731) - partial support
IEEE 802.1AB Link Layer Discovery Protocol
IEEE 802.1D-2004 Spanning Tree Algorithm and traffic priorities
Spanning Tree Protocol
Rapid Spanning Tree Protocol
Multiple Spanning Tree Protocol
IEEE 802.1p Priority tags
IEEE 802.1Q VLAN
IEEE 802.1v Protocol-based VLANs
IEEE 802.1X Port Authentication
IEEE 802.3-2005
Ethernet, Fast Ethernet, Gigabit Ethernet
Link Aggregation Control Protocol (LACP)
Full-duplex flow control (ISO/IEC 8802-3)
IEEE 802.3ac VLAN tagging
IEEE 802.1ag Connectivity Fault Management (Amendment 5, D7.1)
ARP (RFC 826)
DHCP Client (RFC 2131)
DHCP Relay (RFC 951, 2132, 3046)
HTTPS
ICMP (RFC 792)

IGMP (RFC 1112)
IGMPv2 (RFC 2236)
IGMPv3 (RFC 3376) - partial support
IGMP Proxy (RFC 4541)
IPv4 IGMP (RFC 3228)
MLD Snooping (RFC 4541)
NTP (RFC 1305)
RADIUS+ (RFC 2618)
RIPv1 (RFC 1058)
RIPv2 (RFC 2453)
RIPv2, extension (RFC 1724)
RMON (RFC 2819 groups 1,2,3,9)
SNMP (RFC 1157)
SNMPv2c (RFC 1901, 2571)
SNMPv3 (RFC DRAFT 2273, 2576, 3410, 3411, 3413, 3414, 3415)
SNTP (RFC 2030)
SSH (Version 2.0)
TELNET (RFC 854, 855, 856)
TFTP (RFC 1350)

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 (RFC 2054)
Link Aggregation MIB (IEEE 802.3ad)
MAU MIB (RFC 3636)
MIB II (RFC 1213)
P-Bridge MIB (RFC 2674P)
Port Access Entity MIB (IEEE 802.1X)
Port Access Entity Equipment MIB
Power Ethernet MIB (RFC 3621)
Private MIB
Q-Bridge MIB (RFC 2674Q)
QinQ Tunneling (IEEE 802.1ad Provider Bridges)
Quality of Service MIB
RADIUS Accounting Server MIB (RFC 2621)
RADIUS Authentication Client MIB (RFC 2619)
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)



Troubleshooting

Problems Accessing the Management Interface

Table 49: Troubleshooting Chart

Symptom	Action
Cannot connect using a web browser	◆ Be sure the switch is powered on. ◆ Check network cabling between the management station and the switch. Make sure the ends are properly connected and there is no damage to the cable. Test the cable if necessary. ◆ Check that you have a valid network connection to the switch and that the port you are using has not been disabled. ◆ Be sure you have configured the VLAN interface through which the management station is connected with a valid IP address, subnet mask and default gateway. ◆ Be sure the management station has an IP address in the same subnet as the switch's IP interface to which it is connected. ◆ If you are trying to connect to the switch via the IP address for a tagged VLAN group, your management station, and the ports connecting intermediate switches in the network, must be configured with the appropriate tag.
Cannot access the on-board configuration program via a serial port connection	◆ Refer to the <i>CLI Reference Guide</i> for information on troubleshooting a connection to the serial port
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

Glossary

ACL Access Control List. ACLs can limit network traffic and restrict access to certain users or devices by checking each packet for certain IP or MAC (i.e., Layer 2) information.

ARP Address Resolution Protocol converts between IP addresses and MAC (hardware) addresses. ARP is used to locate the MAC address corresponding to a given IP address. This allows the switch to use IP addresses for routing decisions and the corresponding MAC addresses to forward packets from one hop to the next.

BOOTP Boot Protocol is used to provide bootup information for network devices, including IP address information, the address of the TFTP server that contains the devices system files, and the name of the boot file.

CoS Class of Service is supported by prioritizing packets based on the required level of service, and then placing them in the appropriate output queue. Data is transmitted from the queues using weighted round-robin service to enforce priority service and prevent blockage of lower-level queues. Priority may be set according to the port default, the packet's priority bit (in the VLAN tag), TCP/UDP port number, IP Precedence bit, or DSCP priority bit.

DHCP Dynamic Host Control Protocol. Provides a framework for passing configuration information to hosts on a TCP/IP network. DHCP is based on the Bootstrap Protocol (BOOTP), adding the capability of automatic allocation of reusable network addresses and additional configuration options.

DHCP Option 82 A relay option for sending information about the requesting client (or an intermediate relay agent) in the DHCP request packets forwarded by the switch and in reply packets sent back from the DHCP server. This information can be used by DHCP servers to assign fixed IP addresses, or set other services or policies for clients.

DHCP Snooping A technique used to enhance network security by snooping on DHCP server messages to track the physical location of hosts, ensure that hosts only use the IP addresses assigned to them, and ensure that only authorized DHCP servers are accessible.

- DiffServ** Differentiated Services provides quality of service on large networks by employing a well-defined set of building blocks from which a variety of aggregate forwarding behaviors may be built. Each packet carries information (DS byte) used by each hop to give it a particular forwarding treatment, or per-hop behavior, at each network node. DiffServ allocates different levels of service to users on the network with mechanisms such as traffic meters, shapers/droppers, packet markers at the boundaries of the network.
- DNS** Domain Name Service. A system used for translating host names for network nodes into IP addresses.
- DSCP** Differentiated Services Code Point Service. DSCP uses a six-bit tag to provide for up to 64 different forwarding behaviors. Based on network policies, different kinds of traffic can be marked for different kinds of forwarding. The DSCP bits are mapped to the Class of Service categories, and then into the output queues.
- EAPOL** Extensible Authentication Protocol over LAN. EAPOL is a client authentication protocol used by this switch to verify the network access rights for any device that is plugged into the switch. A user name and password is requested by the switch, and then passed to an authentication server (e.g., RADIUS) for verification. EAPOL is implemented as part of the IEEE 802.1X Port Authentication standard.
- EUI** Extended Universal Identifier is an address format used by IPv6 to identify the host portion of the network address. The interface identifier in EUI compatible addresses is based on the link-layer (MAC) address of an interface. Interface identifiers used in global unicast and other IPv6 address types are 64 bits long and may be constructed in the EUI-64 format. The modified EUI-64 format interface ID is derived from a 48-bit link-layer address by inserting the hexadecimal number FFFE between the upper three bytes (OUI field) and the lower 3 bytes (serial number) of the link layer address. To ensure that the chosen address is from a unique Ethernet MAC address, the 7th bit in the high-order byte is set to 1 (equivalent to the IEEE Global/Local bit) to indicate the uniqueness of the 48-bit address.
- GARP** Generic Attribute Registration Protocol. GARP is a protocol that can be used by endstations and switches to register and propagate multicast group membership information in a switched environment so that multicast data frames are propagated only to those parts of a switched LAN containing registered endstations. Formerly called Group Address Registration Protocol.
- GMRP** Generic Multicast Registration Protocol. GMRP allows network devices to register end stations with multicast groups. GMRP requires that any participating network devices or end stations comply with the IEEE 802.1p standard.
- GVRP** GARP VLAN Registration Protocol. Defines a way for switches to exchange VLAN information in order to register necessary VLAN members on ports along the Spanning Tree so that VLANs defined in each switch can work automatically over a Spanning Tree network.

- ICMP** Internet Control Message Protocol is a network layer protocol that reports errors in processing IP packets. ICMP is also used by routers to feed back information about better routing choices.
- IEEE 802.1D** Specifies a general method for the operation of MAC bridges, including the Spanning Tree Protocol.
- IEEE 802.1Q** VLAN Tagging—Defines Ethernet frame tags which carry VLAN information. It allows switches to assign endstations to different virtual LANs, and defines a standard way for VLANs to communicate across switched networks.
- IEEE 802.1p** An IEEE standard for providing quality of service (QoS) in Ethernet networks. The standard uses packet tags that define up to eight traffic classes and allows switches to transmit packets based on the tagged priority value.
- IEEE 802.1s** An IEEE standard for the Multiple Spanning Tree Protocol (MSTP) which provides independent spanning trees for VLAN groups.
- IEEE 802.1w** An IEEE standard for the Rapid Spanning Tree Protocol (RSTP) which reduces the convergence time for network topology changes to about 10% of that required by the older IEEE 802.1D STP standard. (Now incorporated in IEEE 802.1D-2004)
- IEEE 802.1X** Port Authentication controls access to the switch ports by requiring users to first enter a user ID and password for authentication.
- IEEE 802.3ac** Defines frame extensions for VLAN tagging.
- IEEE 802.3x** Defines Ethernet frame start/stop requests and timers used for flow control on full-duplex links. (Now incorporated in IEEE 802.3-2002)
- IGMP** Internet Group Management Protocol. A protocol through which hosts can register with their local router for multicast services. If there is more than one multicast switch/router on a given subnetwork, one of the devices is made the “querier” and assumes responsibility for keeping track of group membership.
- IGMP Proxy** Proxies multicast group membership information onto the upstream interface based on IGMP messages monitored on downstream interfaces, and forwards multicast traffic based on that information. There is no need for multicast routing protocols in a simple tree that uses IGMP Proxy.

IGMP Query On each subnetwork, one IGMP-capable device will act as the querier — that is, the device that asks all hosts to report on the IP multicast groups they wish to join or to which they already belong. The elected querier will be the device with the lowest IP address in the subnetwork.

IGMP Snooping Listening to IGMP Query and IGMP Report packets transferred between IP Multicast Routers and IP Multicast host groups to identify IP Multicast group members.

In-Band Management Management of the network from a station attached directly to the network.

IP Multicast Filtering A process whereby this switch can pass multicast traffic along to participating hosts.

IP Precedence The Type of Service (ToS) octet in the IPv4 header includes three precedence bits defining eight different priority levels ranging from highest priority for network control packets to lowest priority for routine traffic. The eight values are mapped one-to-one to the Class of Service categories by default, but may be configured differently to suit the requirements for specific network applications.

LACP Link Aggregation Control Protocol. Allows ports to automatically negotiate a trunked link with LACP-configured ports on another device.

Layer 2 Data Link layer in the ISO 7-Layer Data Communications Protocol. This is related directly to the hardware interface for network devices and passes on traffic based on MAC addresses.

Layer 3 Network layer in the ISO 7-Layer Data Communications Protocol. This layer handles the routing functions for data moving from one open system to another.

Link Aggregation *See Port Trunk.*

LLDP Link Layer Discovery Protocol is used to discover basic information about neighboring devices in the local broadcast domain by using periodic broadcasts to advertise information such as device identification, capabilities and configuration settings.

MD5 MD5 Message-Digest is an algorithm that is used to create digital signatures. It is intended for use with 32 bit machines and is safer than the MD4 algorithm, which has been broken. MD5 is a one-way hash function, meaning that it takes a message and converts it into a fixed string of digits, also called a message digest.

MIB Management Information Base. An acronym for Management Information Base. It is a set of database objects that contains information about a specific device.

MRD Multicast Router Discovery is a A protocol used by IGMP snooping and multicast routing devices to discover which interfaces are attached to multicast routers. This process allows IGMP-enabled devices to determine where to send multicast source and group membership messages.

MSTP Multiple Spanning Tree Protocol can provide an independent spanning tree for different VLANs. It simplifies network management, provides for even faster convergence than RSTP by limiting the size of each region, and prevents VLAN members from being segmented from the rest of the group.

Multicast Switching A process whereby the switch filters incoming multicast frames for services for which no attached host has registered, or forwards them to all ports contained within the designated multicast VLAN group.

MVR Multicast VLAN Registration is a method of using a single network-wide multicast VLAN to transmit common services, such as such as television channels or video-on-demand, across a service-provider's network. MVR simplifies the configuration of multicast services by using a common VLAN for distribution, while still preserving security and data isolation for subscribers residing in both the MVR VLAN and other standard or private VLAN groups.

NTP Network Time Protocol provides the mechanisms to synchronize time across the network. The time servers operate in a hierarchical-master-slave configuration in order to synchronize local clocks within the subnet and to national time standards via wire or radio.

OSPF Open Shortest Path First is a link-state routing protocol that functions better over a larger network such as the Internet, as opposed to distance-vector routing protocols such as RIP. It includes features such as unlimited hop count, authentication of routing updates, and Variable Length Subnet Masks (VLSM).

Out-of-Band Management Management of the network from a station not attached to the network.

Port Authentication *See IEEE 802.1X.*

Port Mirroring A method whereby data on a target port is mirrored to a monitor port for troubleshooting with a logic analyzer or RMON probe. This allows data on the target port to be studied unobstructively.

Port Trunk Defines a network link aggregation and trunking method which specifies how to create a single high-speed logical link that combines several lower-speed physical links.

Private VLANs Private VLANs provide port-based security and isolation between ports within the assigned VLAN. Data traffic on downlink ports can only be forwarded to, and from, uplink ports.

QinQ QinQ tunneling is designed for service providers carrying traffic for multiple customers across their networks. It is used to maintain customer-specific VLAN and Layer 2 protocol configurations even when different customers use the same internal VLAN IDs.

QoS Quality of Service. QoS refers to the capability of a network to provide better service to selected traffic flows using features such as data prioritization, queuing, congestion avoidance and traffic shaping. These features effectively provide preferential treatment to specific flows either by raising the priority of one flow or limiting the priority of another flow.

RADIUS Remote Authentication Dial-in User Service. RADIUS is a logon authentication protocol that uses software running on a central server to control access to RADIUS-compliant devices on the network.

RIP Routing Information Protocol seeks to find the shortest route to another device by minimizing the distance-vector, or hop count, which serves as a rough estimate of transmission cost. RIP-2 is a compatible upgrade to RIP. It adds useful capabilities for subnet routing, authentication, and multicast transmissions.

RMON Remote Monitoring. RMON provides comprehensive network monitoring capabilities. It eliminates the polling required in standard SNMP, and can set alarms on a variety of traffic conditions, including specific error types.

RSTP Rapid Spanning Tree Protocol. RSTP reduces the convergence time for network topology changes to about 10% of that required by the older IEEE 802.1D STP standard.

SMTP Simple Mail Transfer Protocol is a standard host-to-host mail transport protocol that operates over TCP, port 25.

SNMP Simple Network Management Protocol. The application protocol in the Internet suite of protocols which offers network management services.

- SNTP** Simple Network Time Protocol allows a device to set its internal clock based on periodic updates from a Network Time Protocol (NTP) server. Updates can be requested from a specific NTP server, or can be received via broadcasts sent by NTP servers.
- SSH** Secure Shell is a secure replacement for remote access functions, including Telnet. SSH can authenticate users with a cryptographic key, and encrypt data connections between management clients and the switch.
- STA** Spanning Tree Algorithm is a technology that checks your network for any loops. A loop can often occur in complicated or backup linked network systems. Spanning Tree detects and directs data along the shortest available path, maximizing the performance and efficiency of the network.
- TACACS+** Terminal Access Controller Access Control System Plus. TACACS+ is a logon authentication protocol that uses software running on a central server to control access to TACACS-compliant devices on the network.
- TCP/IP** Transmission Control Protocol/Internet Protocol. Protocol suite that includes TCP as the primary transport protocol, and IP as the network layer protocol.
- Telnet** Defines a remote communication facility for interfacing to a terminal device over TCP/IP.
- TFTP** Trivial File Transfer Protocol. A TCP/IP protocol commonly used for software downloads.
- UDP** User Datagram Protocol. UDP provides a datagram mode for packet-switched communications. It uses IP as the underlying transport mechanism to provide access to IP-like services. UDP packets are delivered just like IP packets – connection-less datagrams that may be discarded before reaching their targets. UDP is useful when TCP would be too complex, too slow, or just unnecessary.
- UTC** Universal Time Coordinate. UTC is a time scale that couples Greenwich Mean Time (based solely on the Earth's rotation rate) with highly accurate atomic time. The UTC does not have daylight saving time.
- VLAN** Virtual LAN. A Virtual LAN is a collection of network nodes that share the same collision domain regardless of their physical location or connection point in the network. A VLAN serves as a logical workgroup with no physical barriers, and allows users to share information and resources as though located on the same LAN.

VRRP Virtual Router Redundancy Protocol uses a virtual IP address to support a primary router and multiple backup routers. The backups can be configured to take over the workload if the master fails or to load share the traffic. The primary goal of VRRP is to allow a host device which has been configured with a fixed gateway to maintain network connectivity in case the primary gateway goes down.

XModem A protocol used to transfer files between devices. Data is grouped in 128-byte blocks and error-corrected.

Index

Numerics

- 802.1Q tunnel 173
 - access 180
 - configuration, guidelines 176
 - configuration, limitations 176
 - CVID to SVID map 178
 - description 173
 - ethernet type 177
 - interface configuration 180
 - mode selection 180
 - status, configuring 177
 - TPID 177
 - uplink 180
- 802.1X
 - authenticator, configuring 349
 - global settings 348
 - port authentication 347
 - port authentication accounting 291, 292

A

- AAA
 - accounting 802.1X port settings 291, 292
 - accounting exec command privileges 292
 - accounting exec settings 291
 - accounting summary 292
 - accounting update 291
 - accounting, configuring 291
 - authorization & accounting 284
 - authorization exec settings 296, 297
 - authorization method 297
 - authorization settings 296
 - authorization summary 298
 - RADIUS group settings 292
 - TACACS+ group settings 292
- acceptable frame type 166
- ACL 324
 - ARP 328, 338
 - binding to a port 340
 - IPv4 Extended 327, 330
 - IPv4 Standard 327, 329
 - IPv6 Extended 327, 334
 - IPv6 Standard 327, 333
 - MAC 328, 336
 - time range 469
- Address Resolution Protocol *See* ARP

- address table 197
 - aging time 199
 - aging time, displaying 199
 - aging time, setting 199
- ARP
 - ACL 338
 - configuration 633
 - description 632
 - proxy 633
 - statistics 636
- ARP inspection 385
 - ACL filter 388
 - additional validation criteria 387
 - ARP ACL 389
 - enabling globally 387
 - trusted ports 389
- ATC 234
 - control response 237
 - functional limitations 235
 - limiting traffic rates 234
 - shutting down a port 235
 - thresholds 237
 - timers 236
 - usage 234
- authentication
 - MAC address authentication 304
 - MAC, configuring ports 307
 - network access 304
 - public key 318
 - web 301
 - web authentication port information, displaying 302
 - web authentication, configuring ports 302
 - web authentication, re-authenticating address 303
 - web authentication, re-authenticating ports 302
 - web, configuring 301
- Automatic Traffic Control *See* ATC

B

- BOOTP 640
- BPDU 206
 - filter 219
 - flooding when STA disabled on VLAN 211
 - flooding when STA globally disabled 211
 - guard 219
 - ignoring superior BPDUs 218

Index

- selecting protocol based on message format 219
- shut down port on receipt 219
- bridge extension capabilities, displaying 75
- broadcast packets, blocking 232
- broadcast storm, threshold 232, 233

C

- cable diagnostics 123
- canonical format indicator 252
- CFM
 - basic operations 507
 - continuity check errors 540
 - continuity check messages 487, 505, 507, 508
 - cross-check message 505, 508
 - cross-check start delay 509
 - delay measure 529
 - description 505
 - domain service access point 505, 517, 521
 - fault isolation 505
 - fault notification 505, 539
 - fault notification generator 508, 514, 539
 - fault verification 505
 - link trace cache 538
 - link trace message 505, 507, 525
 - loop back messages 505, 507, 527
 - maintenance association 505, 517
 - maintenance domain 505, 506, 512
 - maintenance end point 506, 508, 513, 517, 521, 531, 532
 - maintenance intermediate point 506, 513, 534
 - maintenance level 506, 507
 - maintenance point 505
 - MEP archive 515
 - MEP direction 522
 - remote maintenance end point 508, 517, 523, 532, 535, 536
 - service instance 506, 507
 - SNMP traps 510
- class map
 - DiffServ 262
- Class of Service *See* CoS
- clustering switches, management access 464
- committed burst size, QoS policy 269, 270, 271
- committed information rate, QoS policy 269, 270, 271
- community string 429
- configuration files, restoring defaults 77
- configuration settings
 - restoring 79, 80
 - saving 79
- Connectivity Fault Management *See* CFM
- continuity check errors, CFM 540
- continuity check messages, CFM 487, 505, 507, 508
- CoS 241, 249
 - configuring 241

- default mapping to internal values 252
- enabling 248
- layer 3/4 priorities 248
- priorities, mapping to internal values 252
- queue mapping 245
- queue mode 242
- queue weights, assigning 243
- CoS/CFI to PHB/drop precedence 252
- CPU
 - status 98
 - utilization, showing 98
- cross-check message, CFM 505, 508
- cross-check start delay, CFM 509
- CVLAN to SPVLAN map 178

D

- default IPv6 gateway, configuration 643
- default priority, ingress port 241
- default settings, system 42
- delay measure, CFM 529
- DHCP 640, 675
 - class identifier 675
 - client 640
 - client identifier 675
 - relay service 677
 - relay service, enabling 678
- DHCP snooping 356
 - enabling 359
 - global configuration 359
 - information option 359
 - information option policy 359
 - information option, circuit ID 362
 - information option, enabling 359
 - information option, remote ID 359
 - information option, suboption format 359
 - policy selection 359
 - rate limit 359
 - specifying trusted interfaces 361
 - trusted port 362
 - untrusted port 362, 370
 - verifying MAC addresses 359
 - VLAN configuration 361
- DHCPv6 restart 647
- DHCPv6 snooping 364
 - enabling 366
 - global configuration 366
 - specifying trusted interfaces 369
 - trusted port 370
- Differentiated Code Point Service *See* DSCP
- Differentiated Services *See* DiffServ
- DiffServ 261
 - binding policy to interface 275
 - class map 262
 - classifying QoS traffic 262

- color aware, srTCM 270
- color aware, trTCM 271
- color blind, srTCM 270
- color blind, trTCM 271
- committed burst size 270, 271
- committed information rate 270, 271
- configuring 261
- conforming traffic, configuring response 269
- excess burst size 271
- metering, configuring 266, 267
- peak burst size 272
- peak information rate 272
- policy map 265
- policy map, description 262, 269
- QoS policy 265
- service policy 275
- setting CoS for matching packets 269
- setting IP DSCP for matching packets 269
- setting PHB for matching packets 269
- single-rate, three-color meter 266, 270
- srTCM metering 266, 270
- traffic between CIR and BE, configuring response 270
- traffic between CIR and PIR, configuring response 271
- trTCM metering 271
- two-rate, three-color meter 267
- violating traffic, configuring response 272

DNS

- default domain name 669
- displaying the cache 674
- domain name list 669
- enabling lookup 669
- name server list 669
- static entries, IPv4 673
- static entries, IPv6 673

Domain Name Service *See* DNS

domain service access point, CFM 505, 517, 521

downloading software 77

- automatically 81
- using FTP or TFTP 81

drop precedence

- CoS priority mapping 252
- DSCP ingress map 250

DSA encryption 320, 322

DSCP 248, 249

- enabling 248
- mapping to internal values 249

DSCP ingress map, drop precedence 250

DSCP to PHB/drop precedence 250

dynamic addresses

- clearing 198
- displaying 197

dynamic QoS assignment 305, 308

dynamic VLAN assignment 304, 308

E

- edge port, STA 218, 221
- encryption
 - DSA 320, 322
 - RSA 320, 322
- engine ID 431, 432
- ERPS
 - configuration guidelines 474
 - control VLAN 479
 - domain configuration 476
 - domain, enabling 478
 - forced mode 492
 - global configuration 476
 - guard timer 488
 - hold-off timer 488
 - major domain 484
 - manual mode 492
 - MEG level 479
 - neighbor node 480
 - node identifier 484
 - node state 480
 - node type 480
 - non-compliant device protection 487
 - non-revertive, from protection state 480
 - propagate topology change 487
 - protection state, clearing 495
 - R-APS default MAC 486
 - R-APS messages, using MAC 486
 - R-APS with virtual channel 485
 - R-APS, without virtual channel 485
 - revertive recovery 480
 - ring configuration 476
 - ring port, east interface 477
 - ring port, west interface 477
 - ring, enabling 478
 - RPL neighbor 480
 - RPL owner 480
 - secondary ring 484
 - status, displaying 492
 - version 478
 - wait-to-block timer 488
 - wait-to-restore timer 489
 - WTB timer 488
 - WTR timer 489
- event logging 396
- excess burst size, QoS policy 270
- exec command privileges, accounting 292
- exec settings
 - accounting 291
 - authorization 296, 297

F

- fault isolation, CFM 505

Index

fault notification generator, CFM 508, 514, 539
fault notification, CFM 505, 539
fault verification, CFM 505
firmware
 displaying version 73
 upgrading 77
 upgrading automatically 81
 upgrading with FTP or TFTP 81
 version, displaying 73

G

GARP VLAN Registration Protocol See GVRP
gateway, IPv6 default 643
general security measures 283
GNU license 695
GVRP
 enabling 170
 global setting 170
 interface configuration 170

H

hardware version, displaying 73
HTTPS 313, 315
 configuring 313
 replacing SSL certificate 315
 secure-site certificate 315
 UDP port, configuring 314
HTTPS, secure server 313

I

IEEE 802.1D 205
IEEE 802.1s 205
IEEE 802.1w 205
IEEE 802.1X 347
IGMP
 filter profiles, configuration 583, 586
 filter, parameters 583, 586
 filtering & throttling 582
 filtering & throttling, creating profile 584
 filtering & throttling, enabling 583
 filtering & throttling, interface configuration 586
 filtering & throttling, status 583
 groups, displaying 570
 Layer 2 560
 query 562
 snooping 560
 snooping & query, parameters 562
 snooping, configuring 562
 snooping, immediate leave 572
IGMP services, displaying 577
IGMP snooping
 configuring 570
 enabling per interface 570, 572
 filtering query packets 576
 filtering multicast data 576
 forwarding entries 577
 forwarding priority 565
 general query suppression 573
 immediate leave, status 572
 interface attached to multicast router 568
 last leave 562
 last member query count 574
 last member query interval 574
 multicast router discovery 573
 proxy address 574
 proxy reporting 563, 573
 querier status 565
 querier timeout 565
 query interval 574
 query response interval 574
 query suppression 562
 router alert 564
 router port expire time 565
 static host interface 562
 static multicast routing 566
 static port assignment 568
 static router interface 561
 static router port, configuring 566
 statistics, displaying 578
 status 563
 TCN flood 563
 unregistered data flooding 564
 unsolicited report interval 565
 version exclusive 565, 572
 version for interface, setting 573
 version, setting 565
 with proxy reporting 562
immediate leave, IGMP snooping 572
immediate leave, MLD snooping 590
importing user public keys 322
ingress filtering 166
IP address
 BOOTP/DHCP 640
 setting 639
IP filter, for management access 342
IP Port to PHB/drop precedence 258
IP Precedence 249
 enabling 248
IP precedence to PHB/drop precedence 256
IP routing 629, 661
 configuring interfaces 664
 unicast protocols 663
IP source guard
 ACL table, learning mode 374
 learning mode, ACL table or MAC table 374
 MAC table, learning mode 374

- IPv4 address
 - BOOTP/DHCP 640
 - setting 639
 - IPv4 source guard
 - configuring static entries 375
 - setting filter criteria 373
 - setting maximum bindings 374
 - IPv6
 - displaying neighbors 653
 - duplicate address detection 653
 - enabling 645
 - MTU 645
 - ND DAD attempts 645
 - ND NS interval 646
 - ND reachable time 646
 - router advertisements, blocking 647
 - IPv6 address
 - dynamic configuration (global unicast) 650
 - dynamic configuration (link-local) 645
 - EUI format 650
 - EUI-64 setting 650
 - explicit configuration 645
 - global unicast 650
 - link-local 651
 - manual configuration (global unicast) 650
 - manual configuration (link-local) 651
 - setting 643
 - IPv6 source guard
 - configuring static entries 381
 - setting filter criteria 379
 - IPv6 statistics 654
- J**
- jumbo frame 74
- K**
- key
 - private 316
 - public 316
 - user public, importing 322
 - key pair
 - host 316
 - host, generating 320
- L**
- LACP
 - admin key 130
 - configuration 129
 - group attributes, configuring 133
 - group members, configuring 131
 - load balancing 139
 - local parameters 136
 - partner parameters 138
 - protocol message statistics 135
 - protocol parameters 129
 - timeout mode 130
 - timeout, for LACPDU 130
 - last member query interval, IGMP snooping 574
 - LBD 555
 - recover action 556
 - recover time 556
 - transmit interval 556
 - license information, GNU 695
 - Link Layer Discovery Protocol - Media Endpoint Discovery
 - See LLDP-MED
 - Link Layer Discovery Protocol See LLDP
 - link trace cache, CFM 538
 - link trace message, CFM 505, 507, 525
 - link type, STA 217, 221
 - LLDP 401
 - device statistics details, displaying 422
 - device statistics, displaying 421
 - display device information 409, 413
 - displaying remote information 413
 - interface attributes, configuring 403
 - local device information, displaying 409
 - message attributes 403
 - message statistics 421
 - remote information, displaying 420
 - remote port information, displaying 413
 - timing attributes, configuring 401
 - TLV 401, 404
 - TLV, 802.1 405
 - TLV, 802.3 405
 - TLV, basic 404
 - TLV, management address 404
 - TLV, port description 404
 - TLV, system capabilities 404
 - TLV, system description 404
 - TLV, system name 405
 - LLDP-MED 401
 - end-node, extended power-via-MDI 418
 - end-node, inventory 418
 - end-node, location 417
 - end-node, network policy 417
 - notification, status 404
 - TLV 405
 - TLV, civic address 406, 407
 - TLV, extended PoE 406
 - TLV, inventory 406
 - TLV, location 406
 - TLV, MED capabilities 405
 - TLV, network policy 406
 - TLV, PoE 406
 - local engine ID 431

Index

- logging
 - messages, displaying 398
 - syslog traps 399
 - to syslog servers 399
 - log-in, web interface 48
 - logon authentication 299
 - encryption keys 288
 - RADIUS client 287
 - RADIUS server 287
 - sequence 285
 - settings 286
 - TACACS+ client 286
 - TACACS+ server 286
 - logon authentication, settings 288
 - loop back messages, CFM 505, 507, 527
 - loopback detection, STA 208
- ## M
- MAC address authentication 304
 - ports, configuring 307
 - reauthentication 306
 - main menu, web interface 52
 - maintenance association, CFM 505, 517
 - maintenance domain, CFM 505, 506, 512
 - maintenance end point, CFM 506, 508, 513, 517, 521, 531, 532
 - maintenance intermediate point, CFM 506, 513, 534
 - maintenance level, CFM 506, 507
 - maintenance point, CFM 505
 - management access, filtering per address 342
 - management access, IP filter 342
 - Management Information Bases (MIBs) 690
 - matching class settings, classifying QoS traffic 263
 - media-type 107
 - memory
 - status 100
 - utilization, showing 100
 - MEP archive, CFM 515
 - mirror port
 - configuring local traffic 142
 - configuring remote traffic 144
 - mirror port, configuring 142
 - MLD snooping 587
 - configuring 587
 - enabling 588
 - groups, displaying 593, 594
 - immediate leave 590
 - immediate leave, status 590
 - interface attached to multicast router 590, 591
 - max response time 588
 - multicast static router port 590
 - querier 588
 - querier, enabling 588
 - query interval 588
 - query, maximum response time 588
 - robustness value 588
 - static port assignment 592
 - static router port 590
 - unknown multicast, handling 589
 - version 589
 - MRD, multicast router discovery 573
 - MSTP 205, 223
 - global settings, configuring 209, 223
 - global settings, displaying 215
 - interface settings, configuring 216, 227
 - interface settings, displaying 229
 - max hop count 212
 - path cost 228
 - region name 212
 - MTU for IPv6 645
 - multicast filtering 559
 - enabling IGMP snooping 572
 - enabling IGMP snooping per interface 570
 - enabling MLD snooping 588
 - router configuration 566
 - multicast groups 570, 577, 593
 - displaying 570, 577, 593
 - static 568, 570, 592, 593
 - multicast router discovery 571
 - multicast router port, displaying 568, 591
 - multicast services
 - configuring 568, 592
 - displaying 570, 593
 - multicast static router port 566
 - configuring 566
 - configuring for MLD snooping 590
 - multicast storm, threshold 232, 233
 - Multicast VLAN Registration *See* MVR
 - multicast, filtering and throttling 582
 - MVR
 - assigning static multicast groups 605, 621
 - configuring 599
 - description 596
 - domain ID 599
 - forwarding status 604
 - interface status, configuring 603
 - interface status, displaying 605
 - IP for control packets sent upstream 600
 - proxy query interval 614
 - proxy switching 597
 - receiver groups, displaying 607
 - robust value for proxy switching 598
 - running status 599
 - setting interface type 604
 - setting multicast domain 599, 604
 - setting multicast groups 600
 - setting multicast priority 599, 615
 - source port mode 614

- specifying a domain 599, 604
- specifying a VLAN 599
- specifying priority 599
- static binding 605
- static binding, group to port 605
- statistics, displaying 608
- status 599, 604
- using immediate leave 604
- VLAN channel 599

MVR6

- assigning static multicast groups 621
- configuring 615
- domain ID 615
- forwarding status 620
- interface status, configuring 619
- interface status, displaying 621
- IP for control packets sent upstream 615
- proxy switching 613
- receiver groups, displaying 623
- robust value for proxy switching 613
- running status 615, 620
- setting interface type 620
- setting multicast domain 620
- setting multicast groups 616
- specifying a domain 620
- specifying a VLAN 615
- specifying priority 615
- static binding 621
- static binding, group to port 621
- statistics, displaying 624
- using immediate leave 620
- VLAN channel 615

N

network access

- authentication 304
- dynamic QoS assignment 308
- dynamic VLAN assignment 308
- guest VLAN 308
- MAC address filter 308
- port configuration 307
- reauthentication 306
- secure MAC information 312

NTP

- authentication keys, specifying 91
- client, enabling 87
- setting the system clock 89
- specifying servers 89

O

OAM

- active mode 497, 542
- displaying settings and status 496, 541

- enabling on switch ports 496, 541
- event log, displaying 500, 545
- message statistics, displaying 499, 544
- mode selection 497, 542
- passive mode 497, 542
- remote device information, displaying 501, 546
- remote loop back test 502, 547
- setting to active mode 497, 542
- setting to passive mode 497, 542

P

packet block

- broadcast 233
- multicast 232
- unknown multicast 233
- unknown unicast 232, 233

passwords 299

- administrator setting 299

path cost 221

- method 211
- STA 217, 221

peak burst size, QoS policy 271

peak information rate, QoS policy 271

per-hop behavior, DSCP ingress map 250

PoE time range 469

policing traffic, QoS policy 266, 269

policy map

- description 269
- DiffServ 265

port authentication 347

port power

- displaying status 426
- inline 424
- inline status 426
- maximum allocation 424
- priority 425
- showing main power 426

port priority

- configuring 241
- default ingress 241
- STA 216

port security, configuring 344

ports

- autonegotiation 107
- broadcast storm threshold 232, 233
- capabilities 107
- configuring 106
- duplex mode 108
- flow control 108
- forced selection of media type 107
- mirroring 142
- mirroring local traffic 142
- mirroring remote traffic 144
- MTU 108

Index

- multicast storm threshold 232, 233
- speed 108
- statistics 111
- transceiver threshold, auto-set 122
- transceiver threshold, trap 122
- unknown unicast storm threshold 232, 233
- power budgets
 - port 424
 - port priority 425
- power savings
 - configuring 141
 - enabling per port 141
- PPPoE 679–683
- priority, default port ingress 241
- private key 316
- problems, troubleshooting 693
- protocol migration 219
- protocol VLANs 185
 - configuring 186
 - configuring groups 186
 - configuring interfaces 187
 - group configuration 186
 - interface configuration 187
- proxy address, IGMP snooping 574
- proxy ARP 633
- proxy reporting, IGMP snooping 573
- public key 316
- PVID, port native VLAN 166

Q

- QinQ Tunneling *See* 802.1Q tunnel
- QoS 261
 - configuration guidelines 262
 - configuring 261
 - CoS/CFI to PHB/drop precedence 252
 - DSCP to PHB/drop precedence 249
 - dynamic assignment 308
 - IP Port to PHB/drop precedence 258
 - IP precedence to PHB/drop precedence 256
 - matching class settings 263
 - PHB to queue 245
 - selecting CoS, DSCP, IP Precedence 248
- QoS policy
 - committed burst size 269, 270, 271
 - excess burst size 270
 - peak burst size 271
 - policing flow 266, 269
 - srTCM 266
 - srTCM police meter 270
 - trTCM 267
 - trTCM police meter 271
- QoS policy, committed information rate 269, 270, 271
- QoS policy, peak information rate 271
- Quality of Service *See* QoS

- query interval, IGMP snooping 574
- query response interval, IGMP snooping 574
- queue mode, setting 242
- queue weight, assigning to CoS 243

R

- RADIUS
 - logon authentication 287
 - settings 287
- rate limit
 - port 231
 - setting 231
- remote engine ID 432
- remote logging 399
- remote maintenance end point, CFM 508, 517, 523, 532, 535, 536
- restarting the system 101
 - at scheduled times 101
 - showing restart time 104
- RMON 454
 - alarm, displaying settings 456
 - alarm, setting thresholds 454
 - event settings, displaying 459
 - response to alarm setting 457
 - statistics history, collection 459
 - statistics history, displaying 461
 - statistics, collection 462
 - statistics, displaying 463
- root guard 218
- routing table, displaying 666
- RSA encryption 320, 322
- RSTP 205
 - global settings, configuring 209
 - global settings, displaying 215
 - interface settings, configuring 216
 - interface settings, displaying 220

S

- secure shell 316
 - configuration 316
- security, general measures 283
- serial port, configuring 95
- service instance, CFM 506, 507
- sFlow 148
 - configuring receiver 149
 - datagram version 150
 - destination 149
 - maximum datagram 150
 - polling 151
 - receiver socket 150
 - receiver timeout 149
 - sampling 151
- Simple Mail Transfer Protocol *See* SMTP

- Simple Network Management Protocol *See* SNMP
 - single rate three color meter *See* srTCM
 - SMTP
 - event handling 399
 - sending log events 399
 - SNMP 426
 - community string 429
 - enabling traps 446
 - enabling traps, mac-address changes 203
 - filtering IP addresses 342
 - global settings, configuring 428
 - trap manager 446
 - traps, CFM 510
 - users, configuring 441, 443
 - SNMPv3
 - engine ID 431, 432
 - engine identifier, local 431
 - engine identifier, remote 432
 - groups 436
 - local users, configuring 441
 - remote users, configuring 443
 - user configuration 441, 443
 - views 433
 - SNTP
 - setting the system clock 87
 - specifying servers 88
 - software
 - displaying version 73
 - downloading 77
 - version, displaying 73
 - Spanning Tree Protocol *See* STA
 - specifications, software 687
 - srTCM
 - police meter 270
 - QoS policy 266
 - SSH 316
 - authentication retries 319
 - configuring 316
 - downloading public keys for clients 322
 - generating host key pair 320
 - server, configuring 319
 - timeout 319
 - SSL, replacing certificate 315
 - STA 205
 - BPDU filter 219
 - BPDU flooding 211, 216
 - BPDU shutdown 219
 - detecting loopbacks 208
 - edge port 218, 221
 - forward delay 212
 - global settings, configuring 209
 - global settings, displaying 215
 - hello time 211
 - interface settings, configuring 216
 - interface settings, displaying 220
 - link type 217, 221
 - loopback detection 208
 - maximum age 211
 - MSTP interface settings, configuring 227
 - MSTP path cost 228
 - path cost 217, 221
 - path cost method 211
 - port priority 216
 - port/trunk loopback detection 208
 - protocol migration 219
 - transmission limit 211
 - standards, IEEE 689
 - startup files
 - creating 77
 - displaying 77
 - setting 77
 - static addresses, setting 201
 - static routes, configuring 665
 - statistics
 - ARP 636
 - history for port 116
 - history for trunk 116
 - port 111
 - STP 210
 - summary, accounting 292
 - summer time, setting 93
 - switch clustering, for management 464
 - switch settings
 - restoring 79
 - saving 79
 - system clock
 - setting 85
 - setting manually 86
 - setting the time zone 92
 - setting with NTP 89
 - setting with SNTP 87
 - summer time 93
 - system logs 396
 - system software, downloading from server 77
- ## T
- TACACS+
 - logon authentication 286
 - settings 288
 - TCN
 - flood 563
 - general query solicitation 564
 - Telnet
 - configuring 97
 - server, enabling 97
 - time range, ACL 469
 - time range, PoE 469
 - time zone, setting 92

Index

- time, setting 85
- TPID 177
- traffic segmentation 153
 - assigning ports 153
 - enabling 153
 - sessions, assigning ports 155
 - sessions, creating 154
- transceiver data, displaying 120
- transceiver thresholds
 - configuring 121
 - displaying 121
- trap manager 446
- troubleshooting 693
- trTCM
 - police meter 271
 - QoS policy 267
- trunk
 - configuration 125
 - LACP 129
 - static 126
- tunneling unknown VLANs, VLAN trunking 156
- two rate three color meter *See* trTCM
- Type Length Value *See* LLDP TLV

U

- UDLD 550
 - aggressive mode 553
 - configuration 550
 - detection interval 551, 553
 - enable on port 552
 - interface settings 552
 - message interval 553
 - neighbor information 554
 - operation state 553
 - port state 553
 - probe interval 551
 - protocol intervals 551
 - recovery interval 551
 - recovery status 551
- unidirectional link detection 550
- unknown multicast packets, blocking 232
- unknown unicast packets, blocking 232
- unknown unicast storm, threshold 232, 233
- unregistered data flooding, IGMP snooping 564
- upgrading software 77
- user account 299
- user password 299

V

- VLAN trunking 156

- VLANs 159–195
 - 802.1Q tunnel mode 180
 - acceptable frame type 166
 - adding static members 165
 - configuring port members, VLAN index 168
 - creating 163
 - description 159, 195
 - displaying port members 168, 169
 - displaying port members by interface 169
 - displaying port members by interface range 169
 - displaying port members by VLAN index 168
 - dynamic assignment 308
 - egress mode 166
 - ingress filtering 166
 - interface configuration 165
 - IP subnet-based 189
 - MAC-based 191
 - port members, displaying 168, 169
 - protocol 185
 - protocol, configuring 186
 - protocol, configuring groups 186
 - protocol, configuring interfaces 187
 - protocol, group configuration 186
 - protocol, interface configuration 187
 - PVID 166
 - tag swapping 193
 - translation, customer and service provider 193
 - tunneling unknown groups 156
 - voice 277
- voice VLANs 277
 - detecting VoIP devices 278
 - enabling for ports 281
 - identifying client devices 279
- VoIP traffic 277
 - ports, configuring 280
 - telephony OUI, configuring 279
 - voice VLAN, configuring 278
- VoIP, detecting devices 281

W

- web authentication 301
 - address, re-authenticating 303
 - configuring 301
 - configuring ports 302
 - port information, displaying 302, 303
 - ports, configuring 302
 - ports, re-authenticating 302
- web interface
 - access requirements 47
 - configuration buttons 50
 - home page 50
 - menu list 52
 - panel display 51

