



12/28-Port Gigabit Ethernet
Layer 2 Switch

ECS4210-12P

ECS4210-12T

ECS4210-28P

ECS4210-28T

Software Release v1.0.0.24

Web Management Guide

Web Management Guide

ECS4210-12P

Layer 2 Managed PoE Switch with
8 10/100/1000BASE-T (RJ-45) PoE Ports,
2 10/100/1000BASE-T (RJ-45) Ports,
and 2 Gigabit SFP Uplink Ports

ECS4210-12T

Layer 2 Managed Switch with
8 10/100/1000BASE-T (RJ-45) Ports,
and 4 Gigabit SFP Uplink Ports

ECS4210-28P

Layer 2 Managed PoE Switch with
24 10/100/1000BASE-T (RJ-45) PoE Ports,
and 4 Gigabit SFP Uplink Ports

ECS4210-28T

Layer 2 Managed Switch with
24 10/100/1000BASE-T (RJ-45) Ports,
and 4 Gigabit SFP Uplink Ports

ECS4210-12P	149100000219A
ECS4210-12T	149100000241A
ECS4210-28P	149100000217H
ECS4210-28T	149100000217H

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 management interface. An introduction to the switch's key features is also provided. For information on initial configuration, refer to the *CLI Reference Guide*.

The guide includes these sections:

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

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

For information on how to manage the switch through the Command Line Interface (CLI), see the following guide:

CLI Reference Guide

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

Installation Guide

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

Quick Start Guide

Safety and Regulatory Information

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



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



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



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

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

March 2014 Revision

This is the third version of this guide. This guide is valid for software release v1.0.0.24. It includes the following changes.

- ◆ Updated the maximum setting for aging time in "[Changing the Aging Time](#)" on [page 170](#).
- ◆ Added the section "[MLD Snooping \(Snooping and Query for IPv6\)](#)" on [page 494](#).

October 2013 Revision

This is the second version of this guide. This guide is valid for software release v1.0.0.18. It includes the following changes.

- ◆ Added the ECS4210-12T model.
- ◆ Updated parameter section for "[Configuring Port Isolation](#)" on [page 111](#).
- ◆ Updated description of new configuration pages for "[Traffic Segmentation](#)" on [page 128](#).
- ◆ Added mask and priority parameters under "[Configuring MAC-based VLANs](#)" on [page 160](#).
- ◆ Updated description of Action field for Traffic > DiffServ > Add Rule page under "[Creating QoS Policies](#)" on [page 232](#).
- ◆ Updated parameters section under "[Using the Ping Function](#)" on [page 427](#).
- ◆ Updated parameter section and added description of IGMP RADIUS Authentication under "[Filtering Multicast Data at Interfaces](#)" on [page 483](#)

April 2013 Revision

This is the first version of this guide. This guide is valid for software release v1.0.0.12.

Contents

How to Use This Guide	3
Contents	7
Figures	17
Tables	29

Section I	Getting Started	31
	1 Introduction	33
	Key Features	33
	Description of Software Features	34
	System Defaults	39
	2 Using the Web Interface	43
	Connecting to the Web Interface	43
	Navigating the Web Browser Interface	44
	Home Page	44
	Configuration Options	45
	Panel Display	46
	Main Menu	47

Section II	Web Configuration	63
	3 Basic Management Tasks	65
	Displaying System Information	66
	Displaying Hardware/Software Versions	67
	Configuring Support for Jumbo Frames	68
	Displaying Bridge Extension Capabilities	69

Managing System Files	71
Copying Files via FTP/TFTP or HTTP	71
Saving the Running Configuration to a Local File	73
Setting The Start-Up File	74
Showing System Files	74
Automatic Operation Code Upgrade	75
Setting the System Clock	79
Setting the Time Manually	79
Setting the SNTP Polling Interval	80
Configuring NTP	81
Configuring Time Servers	82
Setting the Time Zone	85
Configuring the Console Port	86
Configuring Telnet Settings	88
Displaying CPU Utilization	90
Displaying Memory Utilization	91
Resetting the System	91
4 Interface Configuration	95
Port Configuration	95
Configuring by Port List	95
Configuring by Port Range	97
Displaying Connection Status	98
Configuring Local Port Mirroring	99
Configuring Remote Port Mirroring	101
Showing Port or Trunk Statistics	105
Performing Cable Diagnostics	110
Configuring Port Isolation	111
Trunk Configuration	115
Configuring a Static Trunk	116
Configuring a Dynamic Trunk	118
Displaying LACP Port Counters	123
Displaying LACP Settings and Status for the Local Side	124
Displaying LACP Settings and Status for the Remote Side	126

Traffic Segmentation	128
Enabling Traffic Segmentation	128
Configuring Uplink and Downlink Ports	129
VLAN Trunking	131
5 VLAN Configuration	135
IEEE 802.1Q VLANs	135
Configuring VLAN Groups	138
Adding Static Members to VLANs	140
Configuring Dynamic VLAN Registration	145
IEEE 802.1Q Tunneling	148
Enabling QinQ Tunneling on the Switch	152
Adding an Interface to a QinQ Tunnel	153
Protocol VLANs	154
Configuring Protocol VLAN Groups	155
Mapping Protocol Groups to Interfaces	156
Configuring IP Subnet VLANs	158
Configuring MAC-based VLANs	160
Configuring VLAN Mirroring	162
6 Address Table Settings	165
Setting Static Addresses	165
Configuring MAC Address Isolation	167
Changing the Aging Time	170
Displaying the Dynamic Address Table	171
Clearing the Dynamic Address Table	172
Configuring MAC Address Mirroring	173
Configuring Extended MAC Security	175
7 Spanning Tree Algorithm	179
Overview	179
Configuring Loopback Detection	182
Configuring Global Settings for STA	183
Displaying Global Settings for STA	189
Configuring Interface Settings for STA	190
Displaying Interface Settings for STA	194

Configuring Multiple Spanning Trees	196
Configuring Interface Settings for MSTP	200
8 Congestion Control	203
Rate Limiting	203
Storm Control	204
Automatic Traffic Control	206
Setting the ATC Timers	208
Configuring ATC Thresholds and Responses	209
9 Class of Service	213
Layer 2 Queue Settings	213
Setting the Default Priority for Interfaces	213
Selecting the Queue Mode	214
Mapping CoS Values to Egress Queues	217
Layer 3/4 Priority Settings	220
Setting Priority Processing to DSCP or CoS	220
Mapping Ingress DSCP Values to Internal DSCP Values	221
Mapping CoS Priorities to Internal DSCP Values	224
10 Quality of Service	227
Overview	227
Configuring a Class Map	228
Creating QoS Policies	232
Attaching a Policy Map to a Port	241
11 VoIP Traffic Configuration	243
Overview	243
Configuring VoIP Traffic	244
Configuring Telephony OUI	245
Configuring VoIP Traffic Ports	246
12 Security Measures	249
AAA Authorization and Accounting	250
Configuring Local/Remote Logon Authentication	251
Configuring Remote Logon Authentication Servers	252
Configuring AAA Accounting	257

Configuring AAA Authorization	262
Configuring User Accounts	265
Web Authentication	267
Configuring Global Settings for Web Authentication	267
Configuring Interface Settings for Web Authentication	268
Network Access (MAC Address Authentication)	270
Configuring Global Settings for Network Access	272
Configuring Network Access for Ports	273
Configuring Port Link Detection	275
Configuring a MAC Address Filter	276
Displaying Secure MAC Address Information	278
Configuring HTTPS	279
Configuring Global Settings for HTTPS	279
Replacing the Default Secure-site Certificate	281
Configuring the Secure Shell	282
Configuring the SSH Server	285
Generating the Host Key Pair	286
Importing User Public Keys	288
Access Control Lists	290
Setting A Time Range	291
Showing TCAM Utilization	294
Setting the ACL Name and Type	295
Configuring a Standard IPv4 ACL	297
Configuring an Extended IPv4 ACL	298
Configuring a Standard IPv6 ACL	300
Configuring an Extended IPv6 ACL	302
Configuring a MAC ACL	304
Configuring an ARP ACL	306
Binding a Port to an Access Control List	308
Configuring ACL Mirroring	309
Showing ACL Hardware Counters	311
ARP Inspection	312
Configuring Global Settings for ARP Inspection	313
Configuring VLAN Settings for ARP Inspection	315
Configuring Interface Settings for ARP Inspection	316

Displaying ARP Inspection Statistics	317
Displaying the ARP Inspection Log	318
Filtering IP Addresses for Management Access	319
Configuring Port Security	321
Configuring 802.1X Port Authentication	323
Configuring 802.1X Global Settings	325
Configuring Port Authenticator Settings for 802.1X	326
Configuring Port Supplicant Settings for 802.1X	330
Displaying 802.1X Statistics	332
DoS Protection	335
IP Source Guard	341
Configuring Ports for IP Source Guard	341
Configuring Static Bindings for IP Source Guard	343
Displaying Information for Dynamic IP Source Guard Bindings	344
DHCP Snooping	346
DHCP Snooping Global Configuration	348
DHCP Snooping VLAN Configuration	349
Configuring Ports for DHCP Snooping	350
Displaying DHCP Snooping Binding Information	352
13 Basic Administration Protocols	355
Configuring Event Logging	355
System Log Configuration	355
Remote Log Configuration	358
Link Layer Discovery Protocol	359
Setting LLDP Timing Attributes	360
Configuring LLDP Interface Attributes	361
Configuring LLDP Interface Civic-Address	365
Displaying LLDP Local Device Information	367
Displaying LLDP Remote Device Information	371
Displaying Device Statistics	379
Power over Ethernet	381
Displaying the Switch's Overall PoE Power Budget	382
Setting The Port PoE Power Budget	383

Simple Network Management Protocol	385
Configuring Global Settings for SNMP	387
Setting the Local Engine ID	388
Specifying a Remote Engine ID	389
Setting SNMPv3 Views	390
Configuring SNMPv3 Groups	393
Setting Community Access Strings	398
Configuring Local SNMPv3 Users	399
Configuring Remote SNMPv3 Users	401
Specifying Trap Managers	403
Creating SNMP Notification Logs	407
Showing SNMP Statistics	409
Remote Monitoring	411
Configuring RMON Alarms	412
Configuring RMON Events	414
Configuring RMON History Samples	416
Configuring RMON Statistical Samples	419
Switch Clustering	421
Configuring General Settings for Clusters	422
Cluster Member Configuration	423
Managing Cluster Members	425
14 IP Configuration	427
Using the Ping Function	427
Address Resolution Protocol	429
Setting the ARP Timeout	429
Displaying ARP Entries	430
Setting the Switch's IP Address (IP Version 4)	431
Setting the Switch's IP Address (IP Version 6)	434
Configuring the IPv6 Default Gateway	434
Configuring IPv6 Interface Settings	435
Configuring an IPv6 Address	440
Showing IPv6 Addresses	442
Showing the IPv6 Neighbor Cache	444
Showing IPv6 Statistics	445

Showing the MTU for Responding Destinations	451
15 IP Services	453
Domain Name Service	453
Configuring General DNS Service Parameters	453
Configuring a List of Domain Names	454
Configuring a List of Name Servers	456
Configuring Static DNS Host to Address Entries	457
Displaying the DNS Cache	458
Multicast Domain Name Service	459
Dynamic Host Configuration Protocol	460
Specifying A DHCP Client Identifier	460
Configuring DHCP Relay Option 82	461
16 Multicast Filtering	467
Overview	467
Layer 2 IGMP (Snooping and Query)	468
Configuring IGMP Snooping and Query Parameters	470
Specifying Static Interfaces for a Multicast Router	473
Assigning Interfaces to Multicast Services	476
Setting IGMP Snooping Status per Interface	478
Filtering Multicast Data at Interfaces	483
Displaying Multicast Groups Discovered by IGMP Snooping	484
Displaying IGMP Snooping Statistics	485
Filtering and Throttling IGMP Groups	489
Enabling IGMP Filtering and Throttling	489
Configuring IGMP Filter Profiles	490
Configuring IGMP Filtering and Throttling for Interfaces	492
MLD Snooping (Snooping and Query for IPv6)	494
Configuring MLD Snooping and Query Parameters	494
Setting Immediate Leave Status for MLD Snooping per Interface	496
Specifying Static Interfaces for an IPv6 Multicast Router	496
Assigning Interfaces to IPv6 Multicast Services	498
Showing MLD Snooping Groups and Source List	501
Multicast VLAN Registration	502
Configuring MVR Global Settings	504

Configuring MVR Domain Settings	506
Configuring MVR Group Address Profiles	507
Configuring MVR Interface Status	510
Assigning Static MVR Multicast Groups to Interfaces	512
Displaying MVR Receiver Groups	514
Displaying MVR Statistics	515

Section III	Appendices	521
	A Software Specifications	523
	Software Features	523
	Management Features	524
	Standards	525
	Management Information Bases	525
	B Troubleshooting	527
	Problems Accessing the Management Interface	527
	Using System Logs	528
	C License Information	529
	The GNU General Public License	529
	Glossary	533
	Index	541

Figures

Figure 1: Home Page	44
Figure 2: Front Panel Indicators	46
Figure 3: System Information	67
Figure 4: General Switch Information	68
Figure 5: Configuring Support for Jumbo Frames	69
Figure 6: Displaying Bridge Extension Configuration	70
Figure 7: Copy Firmware	72
Figure 8: Saving the Running Configuration	73
Figure 9: Setting Start-Up Files	74
Figure 10: Displaying System Files	75
Figure 11: Configuring Automatic Code Upgrade	78
Figure 12: Manually Setting the System Clock	80
Figure 13: Setting the Polling Interval for SNTP	81
Figure 14: Configuring NTP	82
Figure 15: Specifying SNTP Time Servers	82
Figure 16: Adding an NTP Time Server	83
Figure 17: Showing the NTP Time Server List	84
Figure 18: Adding an NTP Authentication Key	85
Figure 19: Showing the NTP Authentication Key List	85
Figure 20: Setting the Time Zone	86
Figure 21: Console Port Settings	88
Figure 22: Telnet Connection Settings	89
Figure 23: Displaying CPU Utilization	90
Figure 24: Displaying Memory Utilization	91
Figure 25: Restarting the Switch (Immediately)	93
Figure 26: Restarting the Switch (In)	94
Figure 27: Restarting the Switch (At)	94
Figure 28: Restarting the Switch (Regularly)	94
Figure 29: Configuring Connections by Port List	97

Figure 30: Configuring Connections by Port Range	98
Figure 31: Displaying Port Information	99
Figure 32: Configuring Local Port Mirroring	99
Figure 33: Configuring Local Port Mirroring	100
Figure 34: Displaying Local Port Mirror Sessions	101
Figure 35: Configuring Remote Port Mirroring	101
Figure 36: Configuring Remote Port Mirroring (Source)	104
Figure 37: Configuring Remote Port Mirroring (Intermediate)	105
Figure 38: Configuring Remote Port Mirroring (Destination)	105
Figure 39: Showing Port Statistics (Table)	108
Figure 40: Showing Port Statistics (Chart)	109
Figure 41: Performing Cable Tests	111
Figure 42: Enabling Port Isolation Globally	112
Figure 43: Configuring Port Isolation Profiles	113
Figure 44: Displaying Port Isolation Profiles	113
Figure 45: Assigning Port Isolation Profiles	114
Figure 46: Configuring Static Trunks	116
Figure 47: Creating Static Trunks	117
Figure 48: Adding Static Trunks Members	117
Figure 49: Configuring Connection Parameters for a Static Trunk	118
Figure 50: Showing Information for Static Trunks	118
Figure 51: Configuring Dynamic Trunks	118
Figure 52: Configuring the LACP Aggregator Admin Key	121
Figure 53: Enabling LACP on a Port	121
Figure 54: Configuring LACP Parameters on a Port	122
Figure 55: Configuring Connection Parameters for a Dynamic Trunk	122
Figure 56: Displaying Connection Parameters for Dynamic Trunks	123
Figure 57: Showing Members of Dynamic Trunks	123
Figure 58: Displaying LACP Port Counters	124
Figure 59: Displaying LACP Port Internal Information	126
Figure 60: Displaying LACP Port Remote Information	127
Figure 61: Enabling Traffic Segmentation	129
Figure 62: Configuring Members for Traffic Segmentation	130
Figure 63: Showing Traffic Segmentation Members	131
Figure 64: Configuring VLAN Trunking	131

Figure 65: Configuring VLAN Trunking	133
Figure 66: VLAN Compliant and VLAN Non-compliant Devices	136
Figure 67: Using GVRP	138
Figure 68: Creating Static VLANs	139
Figure 69: Modifying Settings for Static VLANs	140
Figure 70: Showing Static VLANs	140
Figure 71: Configuring Static Members by VLAN Index	143
Figure 72: Configuring Static VLAN Members by Interface	144
Figure 73: Configuring Static VLAN Members by Interface Range	144
Figure 74: Configuring Global Status of GVRP	146
Figure 75: Configuring GVRP for an Interface	147
Figure 76: Showing Dynamic VLANs Registered on the Switch	147
Figure 77: Showing the Members of a Dynamic VLAN	148
Figure 78: QinQ Operational Concept	149
Figure 79: Enabling QinQ Tunneling	153
Figure 80: Adding an Interface to a QinQ Tunnel	154
Figure 81: Configuring Protocol VLANs	156
Figure 82: Displaying Protocol VLANs	156
Figure 83: Assigning Interfaces to Protocol VLANs	158
Figure 84: Showing the Interface to Protocol Group Mapping	158
Figure 85: Configuring IP Subnet VLANs	160
Figure 86: Showing IP Subnet VLANs	160
Figure 87: Configuring MAC-Based VLANs	161
Figure 88: Showing MAC-Based VLANs	162
Figure 89: Configuring VLAN Mirroring	163
Figure 90: Showing the VLANs to Mirror	163
Figure 91: Configuring Static MAC Addresses	166
Figure 92: Displaying Static MAC Addresses	167
Figure 93: Setting Global Status for MAC Address Isolation	169
Figure 94: Configuring MAC Address Isolation	169
Figure 95: Displaying Interfaces Assigned a MAC Address Isolation Mode	170
Figure 96: Setting the Address Aging Time	171
Figure 97: Displaying the Dynamic MAC Address Table	172
Figure 98: Clearing Entries in the Dynamic MAC Address Table	173
Figure 99: Mirroring Packets Based on the Source MAC Address	174

Figure 100: Showing the Source MAC Addresses to Mirror	174
Figure 101: Configuring Extended MAC Security on a VLAN	177
Figure 102: Configuring Extended MAC Security on a Port or Trunk	177
Figure 103: STP Root Ports and Designated Ports	180
Figure 104: MSTP Region, Internal Spanning Tree, Multiple Spanning Tree	181
Figure 105: Common Internal Spanning Tree, Common Spanning Tree, Internal Spanning Tree	181
Figure 106: Configuring Port Loopback Detection	183
Figure 107: Configuring Global Settings for STA (STP)	187
Figure 108: Configuring Global Settings for STA (RSTP)	188
Figure 109: Configuring Global Settings for STA (MSTP)	188
Figure 110: Displaying Global Settings for STA	190
Figure 111: Configuring Interface Settings for STA	193
Figure 112: STA Port Roles	195
Figure 113: Displaying Interface Settings for STA	196
Figure 114: Creating an MST Instance	198
Figure 115: Displaying MST Instances	198
Figure 116: Modifying the Priority for an MST Instance	199
Figure 117: Displaying Global Settings for an MST Instance	199
Figure 118: Adding a VLAN to an MST Instance	200
Figure 119: Displaying Members of an MST Instance	200
Figure 120: Configuring MSTP Interface Settings	202
Figure 121: Displaying MSTP Interface Settings	202
Figure 122: Configuring Rate Limits	204
Figure 123: Configuring Storm Control	206
Figure 124: Storm Control by Limiting the Traffic Rate	206
Figure 125: Storm Control by Shutting Down a Port	207
Figure 126: Configuring ATC Timers	209
Figure 127: Configuring ATC Interface Attributes	211
Figure 128: Setting the Default Port Priority	214
Figure 129: Setting the Queue Mode (Strict)	216
Figure 130: Setting the Queue Mode (WRR)	216
Figure 131: Setting the Queue Mode (Strict and WRR)	217
Figure 132: Mapping CoS Values to Egress Queues	219
Figure 133: Showing CoS Values to Egress Queue Mapping	219

Figure 134: Setting the Trust Mode	221
Figure 135: Configuring DSCP to DSCP Internal Mapping	223
Figure 136: Showing DSCP to DSCP Internal Mapping	223
Figure 137: Configuring CoS to DSCP Internal Mapping	225
Figure 138: Showing CoS to DSCP Internal Mapping	225
Figure 139: Configuring a Class Map	229
Figure 140: Showing Class Maps	230
Figure 141: Adding Rules to a Class Map	231
Figure 142: Showing the Rules for a Class Map	231
Figure 143: Configuring a Policy Map	239
Figure 144: Showing Policy Maps	239
Figure 145: Adding Rules to a Policy Map	240
Figure 146: Showing the Rules for a Policy Map	240
Figure 147: Attaching a Policy Map to a Port	242
Figure 148: Configuring a Voice VLAN	245
Figure 149: Configuring an OUI Telephony List	246
Figure 150: Showing an OUI Telephony List	246
Figure 151: Configuring Port Settings for a Voice VLAN	248
Figure 152: Configuring the Authentication Sequence	252
Figure 153: Authentication Server Operation	252
Figure 154: Configuring Remote Authentication Server (RADIUS)	255
Figure 155: Configuring Remote Authentication Server (TACACS+)	256
Figure 156: Configuring AAA Server Groups	256
Figure 157: Showing AAA Server Groups	257
Figure 158: Configuring Global Settings for AAA Accounting	259
Figure 159: Configuring AAA Accounting Methods	260
Figure 160: Showing AAA Accounting Methods	260
Figure 161: Configuring AAA Accounting Service for 802.1X Service	261
Figure 162: Configuring AAA Accounting Service for Exec Service	261
Figure 163: Displaying a Summary of Applied AAA Accounting Methods	261
Figure 164: Displaying Statistics for AAA Accounting Sessions	262
Figure 165: Configuring AAA Authorization Methods	263
Figure 166: Showing AAA Authorization Methods	264
Figure 167: Configuring AAA Authorization Methods for Exec Service	264
Figure 168: Displaying the Applied AAA Authorization Method	265

Figure 169: Configuring User Accounts	266
Figure 170: Showing User Accounts	267
Figure 171: Configuring Global Settings for Web Authentication	268
Figure 172: Configuring Interface Settings for Web Authentication	269
Figure 173: Configuring Global Settings for Network Access	273
Figure 174: Configuring Interface Settings for Network Access	275
Figure 175: Configuring Link Detection for Network Access	276
Figure 176: Configuring a MAC Address Filter for Network Access	277
Figure 177: Showing the MAC Address Filter Table for Network Access	277
Figure 178: Showing Addresses Authenticated for Network Access	279
Figure 179: Configuring HTTPS	281
Figure 180: Downloading the Secure-Site Certificate	282
Figure 181: Configuring the SSH Server	286
Figure 182: Generating the SSH Host Key Pair	287
Figure 183: Showing the SSH Host Key Pair	288
Figure 184: Copying the SSH User's Public Key	289
Figure 185: Showing the SSH User's Public Key	290
Figure 186: Setting the Name of a Time Range	292
Figure 187: Showing a List of Time Ranges	293
Figure 188: Add a Rule to a Time Range	293
Figure 189: Showing the Rules Configured for a Time Range	294
Figure 190: Showing TCAM Utilization	295
Figure 191: Creating an ACL	296
Figure 192: Showing a List of ACLs	296
Figure 193: Configuring a Standard IPv4 ACL	298
Figure 194: Configuring an Extended IPv4 ACL	300
Figure 195: Configuring a Standard IPv6 ACL	302
Figure 196: Configuring an Extended IPv6 ACL	304
Figure 197: Configuring a MAC ACL	306
Figure 198: Configuring a ARP ACL	308
Figure 199: Binding a Port to an ACL	309
Figure 200: Configuring ACL Mirroring	310
Figure 201: Showing the VLANs to Mirror	310
Figure 202: Showing ACL Statistics	311
Figure 203: Configuring Global Settings for ARP Inspection	314

Figure 204: Configuring VLAN Settings for ARP Inspection	316
Figure 205: Configuring Interface Settings for ARP Inspection	317
Figure 206: Displaying Statistics for ARP Inspection	318
Figure 207: Displaying the ARP Inspection Log	319
Figure 208: Creating an IP Address Filter for Management Access	320
Figure 209: Showing IP Addresses Authorized for Management Access	321
Figure 210: Configuring Port Security	323
Figure 211: Configuring Port Security	324
Figure 212: Configuring Global Settings for 802.1X Port Authentication	326
Figure 213: Configuring Interface Settings for 802.1X Port Authenticator	330
Figure 214: Configuring Interface Settings for 802.1X Port Supplicant	332
Figure 215: Showing Statistics for 802.1X Port Authenticator	334
Figure 216: Showing Statistics for 802.1X Port Supplicant	335
Figure 217: Configuring DoS Protection	340
Figure 218: Setting the Filter Type for IP Source Guard	342
Figure 219: Configuring Static Bindings for IP Source Guard	344
Figure 220: Displaying Static Bindings for IP Source Guard	344
Figure 221: Showing the IP Source Guard Binding Table	345
Figure 222: Configuring Global Settings for DHCP Snooping	349
Figure 223: Configuring DHCP Snooping on a VLAN	350
Figure 224: Configuring the Port Mode for DHCP Snooping	351
Figure 225: Displaying the Binding Table for DHCP Snooping	353
Figure 226: Configuring Settings for System Memory Logs	357
Figure 227: Showing Error Messages Logged to System Memory	358
Figure 228: Configuring Settings for Remote Logging of Error Messages	359
Figure 229: Configuring LLDP Timing Attributes	361
Figure 230: Configuring LLDP Interface Attributes	365
Figure 231: Configuring the Civic Address for an LLDP Interface	366
Figure 232: Showing the Civic Address for an LLDP Interface	367
Figure 233: Displaying Local Device Information for LLDP (General)	370
Figure 234: Displaying Local Device Information for LLDP (Port)	370
Figure 235: Displaying Local Device Information for LLDP (Port Details)	370
Figure 236: Displaying Basic LLDP Information for a Remote Device (Port)	376
Figure 237: Displaying Remote Device Information for LLDP (Port Details)	378
Figure 238: Displaying Remote Device Information for LLDP (End Node)	379

Figure 239: Displaying LLDP Device Statistics (General)	380
Figure 240: Displaying LLDP Device Statistics (Port)	381
Figure 241: Showing the Switch's PoE Budget	382
Figure 242: Setting a Port's PoE Budget	384
Figure 243: Configuring Global Settings for SNMP	388
Figure 244: Configuring the Local Engine ID for SNMP	389
Figure 245: Configuring a Remote Engine ID for SNMP	390
Figure 246: Showing Remote Engine IDs for SNMP	390
Figure 247: Creating an SNMP View	391
Figure 248: Showing SNMP Views	392
Figure 249: Adding an OID Subtree to an SNMP View	392
Figure 250: Showing the OID Subtree Configured for SNMP Views	393
Figure 251: Creating an SNMP Group	397
Figure 252: Showing SNMP Groups	397
Figure 253: Setting Community Access Strings	398
Figure 254: Showing Community Access Strings	399
Figure 255: Configuring Local SNMPv3 Users	400
Figure 256: Showing Local SNMPv3 Users	401
Figure 257: Configuring Remote SNMPv3 Users	402
Figure 258: Showing Remote SNMPv3 Users	403
Figure 259: Configuring Trap Managers (SNMPv1)	406
Figure 260: Configuring Trap Managers (SNMPv2c)	406
Figure 261: Configuring Trap Managers (SNMPv3)	407
Figure 262: Showing Notification Managers	407
Figure 263: Creating SNMP Notification Logs	409
Figure 264: Showing SNMP Notification Logs	409
Figure 265: Showing SNMP Statistics	411
Figure 266: Configuring an RMON Alarm	413
Figure 267: Showing Configured RMON Alarms	414
Figure 268: Configuring an RMON Event	415
Figure 269: Showing Configured RMON Events	416
Figure 270: Configuring an RMON History Sample	417
Figure 271: Showing Configured RMON History Samples	418
Figure 272: Showing Collected RMON History Samples	418
Figure 273: Configuring an RMON Statistical Sample	420

Figure 274: Showing Configured RMON Statistical Samples	420
Figure 275: Showing Collected RMON Statistical Samples	421
Figure 276: Configuring a Switch Cluster	423
Figure 277: Configuring Cluster Members	424
Figure 278: Showing Cluster Members	424
Figure 279: Showing Cluster Candidates	424
Figure 280: Managing a Cluster Member	425
Figure 281: Pinging a Network Device	428
Figure 282: Setting the ARP Timeout	430
Figure 283: Displaying ARP Entries	430
Figure 284: Configuring a Static IPv4 Address	432
Figure 285: Configuring an Auto IP Address	433
Figure 286: Configuring a Dynamic IPv4 Address	433
Figure 287: Configuring the IPv6 Default Gateway	435
Figure 288: Configuring General Settings for an IPv6 Interface	439
Figure 289: Configuring RA Guard for an IPv6 Interface	439
Figure 290: Configuring an IPv6 Address	442
Figure 291: Showing Configured IPv6 Addresses	443
Figure 292: Showing IPv6 Neighbors	445
Figure 293: Showing IPv6 Statistics (IPv6)	449
Figure 294: Showing IPv6 Statistics (ICMPv6)	450
Figure 295: Showing IPv6 Statistics (UDP)	450
Figure 296: Showing Reported MTU Values	451
Figure 297: Configuring General Settings for DNS	454
Figure 298: Configuring a List of Domain Names for DNS	455
Figure 299: Showing the List of Domain Names for DNS	455
Figure 300: Configuring a List of Name Servers for DNS	456
Figure 301: Showing the List of Name Servers for DNS	457
Figure 302: Configuring Static Entries in the DNS Table	457
Figure 303: Showing Static Entries in the DNS Table	458
Figure 304: Showing Entries in the DNS Cache	459
Figure 305: Configuring Multicast DNS	460
Figure 306: Specifying A DHCP Client Identifier	461
Figure 307: Layer 2 DHCP Relay Service	462
Figure 308: Configuring DHCP Relay Information Option 82 Service	465

Figure 309: Multicast Filtering Concept	467
Figure 310: Configuring General Settings for IGMP Snooping	473
Figure 311: Configuring a Static Interface for a Multicast Router	475
Figure 312: Showing Static Interfaces Attached a Multicast Router	475
Figure 313: Showing Current Interfaces Attached a Multicast Router	476
Figure 314: Assigning an Interface to a Multicast Service	477
Figure 315: Showing Static Interfaces Assigned to a Multicast Service	477
Figure 316: Configuring IGMP Snooping on an Interface	482
Figure 317: Showing Interface Settings for IGMP Snooping	483
Figure 318: Dropping IGMP Query Packets	484
Figure 319: Showing Multicast Groups Learned by IGMP Snooping	485
Figure 320: Displaying IGMP Snooping Statistics – Query	487
Figure 321: Displaying IGMP Snooping Statistics – VLAN	488
Figure 322: Displaying IGMP Snooping Statistics – Port	488
Figure 323: Enabling IGMP Filtering and Throttling	490
Figure 324: Creating an IGMP Filtering Profile	491
Figure 325: Showing the IGMP Filtering Profiles Created	491
Figure 326: Adding Multicast Groups to an IGMP Filtering Profile	492
Figure 327: Showing the Groups Assigned to an IGMP Filtering Profile	492
Figure 328: Configuring IGMP Filtering and Throttling Interface Settings	493
Figure 329: Configuring General Settings for MLD Snooping	495
Figure 330: Configuring Immediate Leave for MLD Snooping	496
Figure 331: Configuring a Static Interface for an IPv6 Multicast Router	497
Figure 332: Showing Static Interfaces Attached an IPv6 Multicast Router	498
Figure 333: Showing Current Interfaces Attached an IPv6 Multicast Router	498
Figure 334: Assigning an Interface to an IPv6 Multicast Service	499
Figure 335: Showing Static Interfaces Assigned to an IPv6 Multicast Service	500
Figure 336: Showing Current Interfaces Assigned to an IPv6 Multicast Service	500
Figure 337: Showing IPv6 Multicast Services and Corresponding Sources	502
Figure 338: MVR Concept	503
Figure 339: Configuring Global Settings for MVR	505
Figure 340: Configuring Domain Settings for MVR	507
Figure 341: Configuring an MVR Group Address Profile	508
Figure 342: Displaying MVR Group Address Profiles	508
Figure 343: Assigning an MVR Group Address Profile to a Domain	509

Figure 344: Showing the MVR Group Address Profiles Assigned to a Domain	509
Figure 345: Configuring Interface Settings for MVR	512
Figure 346: Assigning Static MVR Groups to a Port	513
Figure 347: Showing the Static MVR Groups Assigned to a Port	514
Figure 348: Displaying MVR Receiver Groups	515
Figure 349: Displaying MVR Statistics – Query	517
Figure 350: Displaying MVR Statistics – VLAN	518
Figure 351: Displaying MVR Statistics – Port	519

Tables

Table 1: Key Features	33
Table 2: System Defaults	39
Table 3: Web Page Configuration Buttons	45
Table 4: Switch Main Menu	47
Table 5: Port Statistics	106
Table 6: LACP Port Counters	123
Table 7: LACP Internal Configuration Information	124
Table 8: LACP Remote Device Configuration Information	126
Table 9: Traffic Segmentation Forwarding	129
Table 10: MAC Address Isolation Matrix	168
Table 11: Recommended STA Path Cost Range	191
Table 12: Default STA Path Costs	191
Table 13: IEEE 802.1p Egress Queue Priority Mapping	217
Table 14: CoS Priority Levels	218
Table 15: Mapping Internal Per-hop Behavior to Hardware Queues	218
Table 16: Default Mapping of DSCP Values to Internal PHB/Drop Values	222
Table 17: Default Mapping of CoS/CFI to Internal PHB/Drop Precedence	224
Table 18: Dynamic QoS Profiles	271
Table 19: HTTPS System Support	280
Table 20: ARP Inspection Statistics	317
Table 21: ARP Inspection Log	318
Table 22: 802.1X Statistics	332
Table 23: Logging Levels	356
Table 24: LLDP MED Location CA Types	365
Table 25: Chassis ID Subtype	367
Table 26: System Capabilities	368
Table 27: Port ID Subtype	369
Table 28: Remote Port Auto-Negotiation Advertised Capability	372
Table 29: SNMPv3 Security Models and Levels	386

Table 30: Supported Notification Messages	394
Table 31: Address Resolution Protocol	429
Table 32: Show IPv6 Neighbors - display description	444
Table 33: Show IPv6 Statistics - display description	446
Table 34: Show MTU - display description	451
Table 35: RADIUS Server AVPs	484
Table 36: Troubleshooting Chart	527

Section I

Getting Started

This section provides an overview of the switch, and introduces some basic concepts about network switches. It also describes the menu structure for the management interface.

This section includes these chapters:

- ◆ ["Introduction" on page 33](#)
- ◆ ["Using the Web Interface" on page 43](#)

1

Introduction

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

Key Features

Table 1: Key Features

Feature	Description
Configuration Backup and Restore	Using management station or FTP/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 IP Source Guard Port Authentication – IEEE 802.1X Port Isolation - by traffic type and protocol type Port Security – MAC address filtering Sticky Dynamic MAC – Prevents learned address move Traffic Segmentation - by port designation
Access Control Lists	Supports up to 512 rules, 64 ACLs, and a maximum of 32 rules for an ACL
DHCP/DHCPv6	Client
DNS	Client and Proxy service
Port Configuration	Speed and duplex mode and flow control
Port Trunking	Supports up to 8 trunks – static or dynamic trunking (LACP)
Port Mirroring	50 sessions, one or more source ports to one analysis port
Congestion Control	Rate Limiting Throttling for broadcast, multicast, unknown unicast storms
Address Table	16K MAC addresses in the forwarding table, 1K static MAC addresses, 256 L2 multicast groups

Table 1: Key Features (Continued)

Feature	Description
IP Version 4 and 6	Supports IPv4 and IPv6 addressing, and management
IEEE 802.1D Bridge	Supports dynamic data switching and addresses learning
Store-and-Forward Switching	Supported to ensure wire-speed switching while eliminating bad frames
Spanning Tree Algorithm	Supports standard STP, Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Trees (MSTP)
Virtual LANs	Up to 256 using IEEE 802.1Q, port-based, protocol-based, voice VLANs, and QinQ tunnel
Traffic Prioritization	Default port priority, traffic class map, queue scheduling, IP Precedence, or Differentiated Services Code Point (DSCP)
Quality of Service	Supports Differentiated Services (DiffServ)
Link Layer Discovery Protocol	Used to discover basic information about neighboring devices
Multicast Filtering	Supports IGMP snooping and query for Layer 2, MLD snooping and query for Layer 3, and Multicast VLAN Registration
Switch Clustering	Supports up to 36 member switches in a cluster

Description of Software Features

The switch provides a wide range of advanced performance enhancing features. Flow control eliminates the loss of packets due to bottlenecks caused by port saturation. Storm suppression prevents broadcast, multicast, and unknown unicast traffic storms from engulfing the network. Untagged (port-based), tagged, and protocol-based VLANs, plus support for automatic GVRP VLAN registration provide traffic security and efficient use of network bandwidth. CoS priority queueing ensures the minimum delay for moving real-time multimedia data across the network. While multicast filtering 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, and MAC address filtering and IP source guard also provide authenticated for port access. While DHCP snooping is provided to prevent malicious attacks from insecure ports.

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.

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 8 trunks.

STORM CONTROL Broadcast, multicast and unknown unicast storm suppression prevents traffic from overwhelming the network. When enabled on a port, the level of broadcast traffic passing through the port is restricted. If broadcast 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 8 MB for frame buffering. This buffer can queue packets awaiting transmission on congested networks.

SPANNING TREE ALGORITHM The switch supports these spanning tree protocols:

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

convergence than RSTP by limiting the size of each region, and prevents VLAN members from being segmented from the rest of the group (as sometimes occurs with IEEE 802.1D STP).

VIRTUAL LANs The switch supports up to 256 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.
- ◆ 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 four priority queues with strict priority, Weighted Round Robin (WRR) scheduling, or a combination of strict and weighted queuing. It uses IEEE 802.1p and 802.1Q tags to prioritize incoming traffic based on input from the end-station application. These functions can be used to provide independent priorities for delay-sensitive data and best-effort data.

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

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

MULTICAST FILTERING Specific multicast traffic can be assigned to its own VLAN to ensure that it does not interfere with normal network traffic and to guarantee real-time delivery by setting the required priority level for the designated VLAN. The switch uses IGMP Snooping and Query for IPv4 and MLD Snooping and Query for IPv6 to manage multicast group registration. It also supports Multicast VLAN Registration 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	Privileged Exec Level	Username "admin" Password "admin"
	Normal Exec Level	Username "guest" Password "guest"
	Enable Privileged Exec from Normal Exec Level	Password "super"
	RADIUS Authentication	Disabled
	TACACS+ Authentication	Disabled
	802.1X Port Authentication	Disabled
	Web Authentication	Disabled
	MAC Authentication	Disabled
	HTTPS	Enabled
	SSH	Disabled
	Port Security	Disabled
	IP Filtering	Disabled
	DHCP Snooping	Disabled
IP Source Guard	Disabled (all ports)	
Web Management	HTTP Server	Enabled
	HTTP Port Number	80
	HTTP Secure Server	Disabled
	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: Disabled Multicast: Disabled Unknown Unicast: Disabled
Address Table	Aging Time	300 seconds
Spanning Tree Algorithm	Status	Enabled, RSTP (Defaults: RSTP standard)
	Edge Ports	Auto
LLDP	Status	Enabled
Virtual LANs	Default VLAN	1
	PVID	1
	Acceptable Frame Type	All
	Ingress Filtering	Disabled
	Switchport Mode (Egress Mode)	Access
	GVRP (global)	Disabled
	GVRP (port interface)	Disabled
	QinQ Tunneling	Disabled
Traffic Prioritization	Ingress Port Priority	0
	Queue Mode	Strict-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

Table 2: System Defaults (Continued)

Function	Parameter	Default
IP Settings	Management. VLAN	VLAN 1
	IP Address	DHCP
	Subnet Mask	255.255.0.0
	Default Gateway	0.0.0.0
	DHCP	Client: Disabled
	DNS	Proxy service
	Multicast DNS	Enabled
	BOOTP	Disabled
Multicast Filtering	IGMP Snooping (Layer 2)	Snooping: Disabled Querier: Disabled
	MLD Snooping (Layer 2 IPv6)	Snooping: Enabled Querier: Disabled
	IGMP Proxy Reporting	Disabled
System Log	Status	Enabled
	Messages Logged to RAM	Levels 0-7 (all)
	Messages Logged to Flash	Levels 0-3
SNTP	Clock Synchronization	Disabled
Switch Clustering	Status	Disabled
	Commander	Disabled

2

Using the Web Interface

This switch provides an embedded HTTP web agent. Using a web browser you can configure the switch and view statistics to monitor network activity. The web agent can be accessed by any computer on the network using a standard web browser (Internet Explorer 6.x or above, or Mozilla Firefox 4.x or above).



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 “Using the Command Line Interface” in 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 “Setting an IP Address” 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 “Setting Passwords” in 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 190.

Note: Users are automatically logged off of the HTTP server or HTTPS server if no input is detected for 300 seconds.

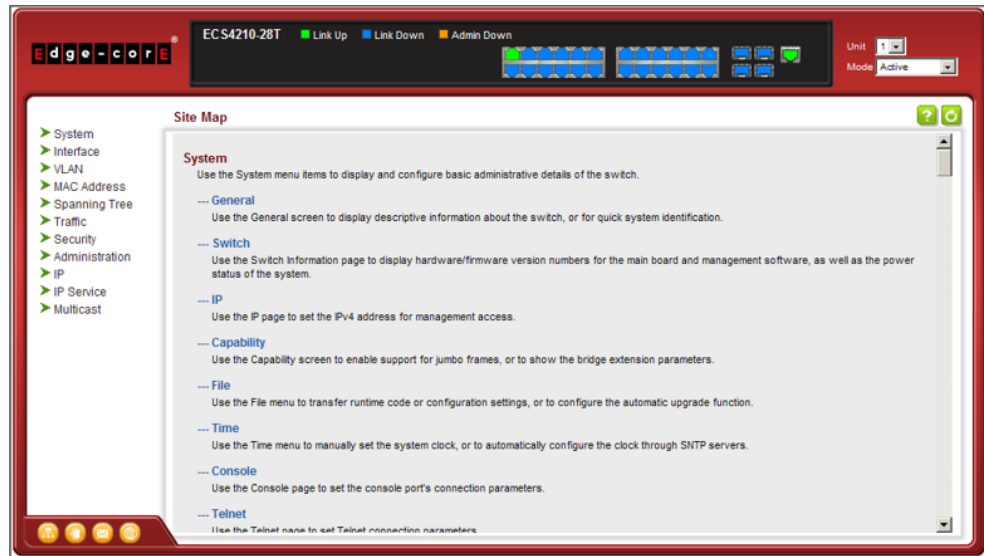
Note: Connection to the web interface is not supported for HTTPS using an IPv6 link local address.

Navigating the Web Browser Interface

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

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

Figure 1: Home Page










Note: This manual covers the ECS4210-12P and ECS4210-28P Gigabit Ethernet PoE switches, as well as the ECS4210-12T and ECS4210-28T Gigabit Ethernet switches. Other than the number of ports and support for PoE, there are no other significant differences. Therefore nearly all of the screen display examples are based on the ECS4210-28T. The panel graphics for all 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.”
	Save current configuration settings.
	Displays help for the selected page.
	Refreshes the current page.
	Displays the site map.
	Logs out of the management interface.
	Links to the manufacture’s web site.
	Sends mail to the manufacturer.


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

Figure 2: Front Panel Indicators


ECS4210-12P




ECS4210-28P



ECS4210-12T



ECS4210-28T



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

Table 4: Switch Main Menu

Menu	Description	Page
System		
General	Provides basic system description, including contact information	66
Switch	Shows the number of ports, hardware version, power status, and firmware version numbers	67
IP	Sets the IPv4 address for management access	431
Capability	Enables support for jumbo frames; shows the bridge extension parameters	68, 69
File		71
Copy	Allows the transfer and copying files	71
Set Startup	Sets the startup file	74
Show	Shows the files stored in flash memory; allows deletion of files	74
Automatic Operation Code Upgrade	Automatically upgrades operation code if a newer version is found on the server	75
Time		79
Configure General		
Manual	Manually sets the current time	79
SNTP	Configures SNTP polling interval	80
NTP	Configures NTP authentication parameters	81
Configure Time Server	Configures a list of SNTP servers	82
Configure SNTP Server	Sets the IP address for SNTP time servers	82
Add NTP Server	Adds NTP time server and index of authentication key	83
Show NTP Server	Shows list of configured NTP time servers	83
Add NTP Authentication Key	Adds key index and corresponding MD5 key	84
Show NTP Authentication Key	Shows list of configured authentication keys	84
Configure Time Zone	Sets the local time zone for the system clock	85
Console	Sets console port connection parameters	86
Telnet	Sets Telnet connection parameters	88
CPU Utilization	Displays information on CPU utilization	90
Memory Status	Shows memory utilization parameters	91

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Reset	Restarts the switch immediately, at a specified time, after a specified delay, or at a periodic interval	91
Interface		95
Port		95
General		
Configure by Port List	Configures connection settings per port	95
Configure by Port Range	Configures connection settings for a range of ports	97
Show Information	Displays port connection status	98
Mirror		99
Add	Sets the source and target ports for mirroring	99
Show	Shows the configured mirror sessions	99
Statistics	Shows Interface, Etherlike, and RMON port statistics	105
Chart	Shows Interface, Etherlike, and RMON port statistics	105
Cable Test	Performs cable diagnostics for selected port to diagnose any cable faults (short, open etc.) and report the cable length	110
Isolation	Restricts the traffic types or protocol types allowed to pass between specified ports	111
Trunk		
Static		116
Configure Trunk	Creates a trunk, specifying port members	116
Configure General		116
Configure	Configures trunk connection settings	116
Show Information	Displays trunk connection settings	116
Dynamic		118
Configure Aggregator	Configures administration key and timeout for LACP groups	118
Configure Aggregation Port		118
Configure		118
General	Allows ports to dynamically join trunks	118
Actor	Configures parameters for link aggregation group members on the local side	118
Partner	Configures parameters for link aggregation group members on the remote side	118
Show Information		123
Counters	Displays statistics for LACP protocol messages	123

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Internal	Displays configuration settings and operational state for the local side of a link aggregation	124
Neighbors	Displays configuration settings and operational state for the remote side of a link aggregation	126
Configure Trunk		118
Show	Displays trunk connection settings	118
Configure	Configures trunk connection settings	118
Show Member	Show port members of dynamic trunks	118
Statistics	Shows Interface, Etherlike, and RMON port statistics	105
Chart	Shows Interface, Etherlike, and RMON port statistics	105
RSPAN	Mirrors traffic from remote switches for analysis at a destination port on the local switch	101
Traffic Segmentation		128
Configure Global	Enables traffic segmentation globally	128
Configure Session	Configures the uplink and down-link ports for a segmented group of ports	129
VLAN Trunking	Allows unknown VLAN groups to pass through the specified interface	131
VLAN	Virtual LAN	135
Static		
Add	Configures VLAN groups, administrative status, and remote type	138
Show	Displays configured VLAN groups	138
Modify	Configures group name and administrative status	138
Edit Member by VLAN	Specifies VLAN attributes per VLAN	140
Edit Member by Interface	Specifies VLAN attributes per interface	140
Edit Member by Interface Range	Specifies VLAN attributes per interface range	140
Dynamic		
Configure General	Enables GVRP VLAN registration protocol globally	145
Configure Interface	Configures GVRP status and timers per interface	145
Show Dynamic VLAN		145
Show VLAN	Shows the VLANs this switch has joined through GVRP	145
Show VLAN Member	Shows the interfaces assigned to a VLAN through GVRP	145
Tunnel	IEEE 802.1Q (QinQ) Tunneling	148
Configure Global	Sets tunnel mode for the switch	152

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Configure Interface	Sets the tunnel mode for any participating interface	153
Protocol		154
Configure Protocol		155
Add	Creates a protocol group, specifying supported protocols	155
Show	Shows configured protocol groups	155
Configure Interface		156
Add	Maps a protocol group to a VLAN	156
Show	Shows the protocol groups mapped to each VLAN	156
IP Subnet		158
Add	Maps IP subnet traffic to a VLAN	158
Show	Shows IP subnet to VLAN mapping	158
MAC-Based		160
Add	Maps traffic with specified source MAC address to a VLAN	160
Show	Shows source MAC address to VLAN mapping	160
Mirror		162
Add	Mirrors traffic from one or more source VLANs to a target port	162
Show	Shows mirror list	162
MAC Address		165
Static		165
Configure Global	Enable MAC address isolation globally on the switch	167
Configure MAC Address		165
Add	Configures static entries in the address table	165
Show	Displays static entries in the address table	165
Dynamic		
Configure Aging	Sets timeout for dynamically learned entries	170
Show Dynamic MAC	Displays dynamic entries in the address table	171
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	172
Mirror		173
Add	Mirrors traffic matching a specified source address from any port on the switch to a target port	173
Show	Shows mirror list	173

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Extended MAC Security	Configures the maximum number of MAC addresses that can be learned on an interface, the movable-static function which allows a static address to be moved to another interface, and the sticky-dynamic function which prevents dynamic address already learned elsewhere from being learned at a specified interface.	175
Spanning Tree		179
Loopback Detection	Configures Loopback Detection parameters	182
STA	Spanning Tree Algorithm	
Configure Global		
Configure	Configures global bridge settings for STP, RSTP and MSTP	183
Show Information	Displays STA values used for the bridge	189
Configure Interface		
Configure	Configures interface settings for STA	190
Show Information	Displays interface settings for STA	194
MSTP	Multiple Spanning Tree Algorithm	196
Configure Global		196
Add	Configures initial VLAN and priority for an MST instance	196
Show	Shows configured MST instances	196
Modify	Modifies priority for an MST instance	196
Add Member	Adds VLAN members for an MST instance	196
Show Member	Adds or deletes VLAN members for an MST instance	196
Show Information	Shows global settings for an MST instance	196
Configure Interface		200
Configure	Configures interface settings for an MST instance	200
Show Information	Displays interface settings for an MST instance	200
Traffic		
Rate Limit	Sets the input and output rate limits for a port	203
Storm Control	Sets the traffic storm threshold for each interface	204
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	204
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	208
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	209

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Priority		213
Default Priority	Sets the default priority for each port or trunk	213
Queue	Sets queue mode for the switch; sets the service weight for each queue that will use a weighted or hybrid mode	214
Trust Mode	Selects IP Precedence, DSCP or CoS priority processing	220
DSCP to DSCP		221
Add	Maps DSCP values in incoming packets to per-hop behavior and drop precedence values for internal priority processing	221
Show	Shows the DSCP to DSCP mapping list	221
CoS to DSCP		224
Add	Maps CoS/CFI values in incoming packets to per-hop behavior and drop precedence values for priority processing	224
Show	Shows the CoS to DSCP mapping list	224
PHB to Queue		217
Add	Maps internal per-hop behavior values to hardware queues	217
Show	Shows the PHB to Queue mapping list	217
DiffServ		227
Configure Class		228
Add	Creates a class map for a type of traffic	228
Show	Shows configured class maps	228
Modify	Modifies the name of a class map	228
Add Rule	Configures the criteria used to classify ingress traffic	228
Show Rule	Shows the traffic classification rules for a class map	228
Configure Policy		232
Add	Creates a policy map to apply to multiple interfaces	232
Show	Shows configured policy maps	232
Modify	Modifies the name of a policy map	232
Add Rule	Sets the boundary parameters used for monitoring inbound traffic, and the action to take for conforming and non-conforming traffic	232
Show Rule	Shows the rules used to enforce bandwidth policing for a policy map	232
Configure Interface	Applies a policy map to an ingress port	241
VoIP	Voice over IP	243
Configure Global	Configures auto-detection of VoIP traffic, sets the Voice VLAN, and VLAN aging time	244

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Configure OUI		245
Add	Maps the OUI in the source MAC address of ingress packets to the VoIP device manufacturer	245
Show	Shows the OUI telephony list	245
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	246
Security		249
AAA	Authentication, Authorization and Accounting	250
System Authentication	Configures authentication sequence – local, RADIUS, and TACACS	251
Server		252
Configure Server	Configures RADIUS and TACACS server message exchange settings	252
Configure Group		252
Add	Specifies a group of authentication servers and sets the priority sequence	252
Show	Shows the authentication server groups and priority sequence	252
Accounting	Enables accounting of requested services for billing or security purposes	257
Configure Global	Specifies the interval at which the local accounting service updates information to the accounting server	257
Configure Method		257
Add	Configures accounting for various service types	257
Show	Shows the accounting settings used for various service types	257
Configure Service	Sets the accounting method applied to specific interfaces for 802.1X, CLI command privilege levels for the console port, and for Telnet	257
Show Information		257
Summary	Shows the configured accounting methods, and the methods applied to specific interfaces	257
Statistics	Shows basic accounting information recorded for user sessions	257
Authorization	Enables authorization of requested services	262
Configure Method		262
Add	Configures authorization for various service types	262
Show	Shows the authorization settings used for various service types	262
Configure Service	Sets the authorization method applied used for the console port, and for Telnet	262
Show Information	Shows the configured authorization methods, and the methods applied to specific interfaces	262

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
User Accounts		265
Add	Configures user names, passwords, and access levels	265
Show	Shows authorized users	265
Modify	Modifies user attributes	265
Web Authentication	Allows authentication and access to the network when 802.1X or Network Access authentication are infeasible or impractical	267
Configure Global	Configures general protocol settings	267
Configure Interface	Enables Web Authentication for individual ports	268
Network Access	MAC address-based network access authentication	270
Configure Global	Enables aging for authenticated MAC addresses, and sets the time period after which a connected MAC address must be reauthenticated	272
Configure Interface		273
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	273
Link Detection	Configures detection of changes in link status, and the response (i.e., send trap or shut down port)	275
Configure MAC Filter		276
Add	Specifies MAC addresses exempt from authentication	276
Show	Shows the list of exempt MAC addresses	276
Show Information	Shows the authenticated MAC address list	278
HTTPS	Secure HTTP	279
Configure Global	Enables HTTPSs, and specifies the UDP port to use	279
Copy Certificate	Replaces the default secure-site certificate	281
SSH	Secure Shell	282
Configure Global	Configures SSH server settings	285
Configure Host Key		286
Generate	Generates the host key pair (public and private)	286
Clear	Displays RSA and DSA host keys; deletes host keys	286
Configure User Key		288
Copy	Imports user public keys from TFTP server	288
Show	Displays RSA and DSA user keys; deletes user keys	288

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
ACL	Access Control Lists	290
Configure ACL		
Show TCAM	Shows utilization parameters for TCAM	294
Add	Adds an ACL based on IP or MAC address filtering	295
Show	Shows the name and type of configured ACLs	295
Add Rule	Configures packet filtering based on IP or MAC addresses and other packet attributes	295
Show Rule	Shows the rules specified for an ACL	295
Configure Interface	Binds a port to the specified ACL and time range	308
ARP Inspection		312
Configure General	Enables inspection globally, configures validation of additional address components, and sets the log rate for packet inspection	313
Configure VLAN	Enables ARP inspection on specified VLANs	315
Configure Interface	Sets the trust mode for ports, and sets the rate limit for packet inspection	316
Show Information		
Show Statistics	Displays statistics on the inspection process	317
Show Log	Shows the inspection log list	318
IP Filter		319
Add	Sets IP addresses of clients allowed management access via the web, SNMP, and Telnet	319
Show	Shows the addresses to be allowed management access	319
Port Security	Configures per port security, including status, response for security breach, and maximum allowed MAC addresses	321
Port Authentication	IEEE 802.1X	323
Configure Global	Enables authentication and EAPOL pass-through	325
Configure Interface	Sets authentication parameters for individual ports	
Authenticator	Sets port authenticator settings	326
Supplicant	Sets port supplicant settings	330
Show Statistics	Displays protocol statistics for the selected port	332
Authenticator	Displays protocol statistics for port authenticator	332
Supplicant	Displays protocol statistics for port supplicant	332
DoS Protection	Protects against Denial-of-Service attacks	335

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
IP Source Guard	Filters IP traffic based on static entries in the IP Source Guard table, or dynamic entries in the DHCP Snooping table	341
Port Configuration	Enables IP source guard and selects filter type per port	341
Static Binding		343
Add	Adds a static addresses to the source-guard binding table	343
Show	Shows static addresses in the source-guard binding table	343
Dynamic Binding	Displays the source-guard binding table for a selected interface	344
Administration		355
Log		355
System		355
Configure Global	Stores error messages in local memory	355
Show Logs	Shows logged error messages	355
Remote	Configures the logging of messages to a remote logging process	358
LLDP		359
Configure Global	Configures global LLDP timing parameters	360
Configure Interface		
Configure General	Sets the message transmission mode, enables SNMP notification, and sets the LLDP attributes to advertise	361
Add CA-Type	Specifies the location of the device attached to an interface	365
Show CA-Type	Shows the location of the device attached to an interface	365
Modify CA-Type	Modifies the location of the device attached to an interface	365
Show Local Device Information		367
General	Displays general information about the local device	367
Port/Trunk	Displays information about each interface	367
Show Remote Device Information		371
Port/Trunk	Displays information about a remote device connected to a port on this switch	371
Port/Trunk Details	Displays detailed information about a remote device connected to this switch	371
Show Device Statistics		379
General	Displays statistics for all connected remote devices	379
Port/Trunk	Displays statistics for remote devices on a selected port or trunk	379

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
PoE*	Power over Ethernet	381
Configure Global	Displays the power budget for the switch	382
Configure Interface	Configures port power parameters	383
SNMP	Simple Network Management Protocol	385
Configure Global	Enables SNMP agent status, and sets related trap functions	387
Configure Engine		388
Set Engine ID	Sets the SNMP v3 engine ID on this switch	388
Add Remote Engine	Sets the SNMP v3 engine ID for a remote device	389
Show Remote Engine	Shows configured engine ID for remote devices	389
Configure View		390
Add View	Adds an SNMP v3 view of the OID MIB	390
Show View	Shows configured SNMP v3 views	390
Add OID Subtree	Specifies a part of the subtree for the selected view	390
Show OID Subtree	Shows the subtrees assigned to each view	390
Configure Group		393
Add	Adds a group with access policies for assigned users	393
Show	Shows configured groups and access policies	393
Configure User		
Add Community	Configures community strings and access mode	398
Show Community	Shows community strings and access mode	398
Add SNMPv3 Local User	Configures SNMPv3 users on this switch	399
Show SNMPv3 Local User	Shows SNMPv3 users configured on this switch	399
Change SNMPv3 Local User Group	Assign a local user to a new group	399
Add SNMPv3 Remote User	Configures SNMPv3 users from a remote device	401
Show SNMPv3 Remote User	Shows SNMPv3 users set from a remote device	401
Configure Trap		403
Add	Configures trap managers to receive messages on key events that occur this switch	403
Show	Shows configured trap managers	403

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
RMON	Remote Monitoring	411
Configure Global		
Add		
Alarm	Sets threshold bounds for a monitored variable	412
Event	Creates a response event for an alarm	414
Show		
Alarm	Shows all configured alarms	412
Event	Shows all configured events	414
Configure Interface		
Add		
History	Periodically samples statistics on a physical interface	416
Statistics	Enables collection of statistics on a physical interface	419
Show		
History	Shows sampling parameters for each entry in the history group	416
Statistics	Shows sampling parameters for each entry in the statistics group	419
Show Details		
History	Shows sampled data for each entry in the history group	416
Statistics	Shows sampled data for each entry in the history group	419
Cluster		421
Configure Global	Globally enables clustering for the switch; sets Commander status	422
Configure Member		
Add	Adds switch Members to the cluster	423
Show Candidate	Shows cluster candidates	423
Show Member	Shows cluster switch member; managed switch members	425
IP		427
General		
Ping	Sends ICMP echo request packets to another node on the network	427
ARP	Address Resolution Protocol	429
Configure General	Sets the aging time for dynamic entries in the ARP cache	429
Show Information	Shows entries in the Address Resolution Protocol (ARP) cache	430
IPv6 Configuration		434
Configure Global	Sets an IPv6 default gateway for traffic with no known next hop	434

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Configure Interface	Configures IPv6 interface address using auto-configuration or link-local address, and sets related protocol settings	435
Add IPv6 Address	Adds an global unicast, EUI-64, or link-local IPv6 address to an interface	440
Show IPv6 Address	Show the IPv6 addresses assigned to an interface	442
Show IPv6 Neighbor Cache	Displays information in the IPv6 neighbor discovery cache	444
Show Statistics		445
IPv6	Shows statistics about IPv6 traffic	445
ICMPv6	Shows statistics about ICMPv6 messages	445
UDP	Shows statistics about UDP messages	445
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	451
IP Service		453
DNS	Domain Name Service	
General		453
Configure Global	Enables DNS lookup; defines the default domain name appended to incomplete host names	453
Add Domain Name	Defines a list of domain names that can be appended to incomplete host names	454
Show Domain Names	Shows the configured domain name list	454
Add Name Server	Specifies IP address of name servers for dynamic lookup	456
Show Name Servers	Shows the name server address list	456
Static Host Table		457
Add	Configures static entries for domain name to address mapping	457
Show	Shows the list of static mapping entries	457
Modify	Modifies the static address mapped to the selected host name	457
Cache	Displays cache entries discovered by designated name servers	458
Multicast DNS	Configures multicast DNS host name-to-address mapping on the local network without the need for a dedicated DNS server	459
DHCP	Dynamic Host Configuration Protocol	
Client	Specifies the DHCP client identifier for an interface	460
Relay Option 82	Specifies DHCP relay servers, including DHCP option 82 information	461
Snooping		346
Configure Global	Enables DHCP snooping globally, MAC-address verification, information option; and sets the information policy	348

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Configure VLAN	Enables DHCP snooping on a VLAN	349
Configure Interface	Sets the trust mode for an interface	350
Show Information	Displays the DHCP Snooping binding information	352
Multicast		467
IGMP Snooping		468
General	Enables multicast filtering; configures parameters for multicast snooping	470
Multicast Router		473
Add Static Multicast Router	Assigns ports that are attached to a neighboring multicast router	473
Show Static Multicast Router	Displays ports statically configured as attached to a neighboring multicast router	473
Show Current Multicast Router	Displays ports attached to a neighboring multicast router, either through static or dynamic configuration	473
IGMP Member		476
Add Static Member	Statically assigns multicast addresses to the selected VLAN	476
Show Static Member	Shows multicast addresses statically configured on the selected VLAN	476
Interface		478
Configure VLAN	Configures IGMP snooping per VLAN interface	478
Show VLAN Information	Shows IGMP snooping settings per VLAN interface	478
Configure Port	Configures the interface to drop IGMP query packets	483
Configure Trunk	Configures the interface to drop IGMP query packets	483
Forwarding Entry	Displays the current multicast groups learned through IGMP Snooping	484
Filter		489
Configure General	Enables IGMP filtering for the switch	489
Configure Profile		490
Add	Adds IGMP filter profile; and sets access mode	490
Show	Shows configured IGMP filter profiles	490
Add Multicast Group Range	Assigns multicast groups to selected profile	490
Show Multicast Group Range	Shows multicast groups assigned to a profile	490
Configure Interface	Assigns IGMP filter profiles to port interfaces and sets throttling action	492
Statistics		485
Show Query Statistics	Shows statistics for query-related messages	485
Show VLAN Statistics	Shows statistics for protocol messages, number of active groups	485

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show Port Statistics	Shows statistics for protocol messages, number of active groups	485
Show Trunk Statistics	Shows statistics for protocol messages, number of active groups	485
MLD Snooping		494
General	Enables multicast filtering; configures parameters for IPv6 multicast snooping	494
Interface	Configures Immediate Leave status for a VLAN	496
Multicast Router		496
Add Static Multicast Router	Assigns ports that are attached to a neighboring multicast router	496
Show Static Multicast Router	Displays ports statically configured as attached to a neighboring multicast router	496
Show Current Multicast Router	Displays ports attached to a neighboring multicast router, either through static or dynamic configuration	496
MLD Member		498
Add Static Member	Statically assigns multicast addresses to the selected VLAN	498
Show Static Member	Shows multicast addresses statically configured on the selected VLAN	498
Show Current Member	Shows multicast addresses associated with the selected VLAN, either through static or dynamic configuration	498
Group Information	Displays known multicast groups, member ports, the means by which each group was learned, and the corresponding source list	498
MVR	Multicast VLAN Registration	502
Configure Global	Configures proxy switching and robustness value	504
Configure Domain	Enables MVR for a domain, sets the MVR VLAN, forwarding priority, and upstream source IP	506
Configure Profile		507
Add	Configures multicast stream addresses	507
Show	Shows multicast stream addresses	507
Associate Profile		507
Add	Maps an address profile to a domain	507
Show	Shows addresses profile to domain mapping	507
Configure Interface	Configures MVR interface type and immediate leave mode; also displays MVR operational and active status	510
Configure Static Group Member		512
Add	Statically assigns MVR multicast streams to an interface	512
Show	Shows MVR multicast streams assigned to an interface	512
Show Member	Shows the multicast groups assigned to an MVR VLAN, and the source address of the multicast services	514

Table 4: Switch Main Menu (Continued)

Menu	Description	Page
Show Statistics		515
Show Query Statistics	Shows statistics for query-related messages	515
Show VLAN Statistics	Shows statistics for protocol messages and number of active groups	515
Show Port Statistics	Shows statistics for protocol messages and number of active groups	515
Show Trunk Statistics	Shows statistics for protocol messages and number of active groups	515

* ECS4210-12P, ECS4210-28P

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:

- ◆ [“Basic Management Tasks” on page 65](#)
- ◆ [“Interface Configuration” on page 95](#)
- ◆ [“VLAN Configuration” on page 135](#)
- ◆ [“Address Table Settings” on page 165](#)
- ◆ [“Spanning Tree Algorithm” on page 179](#)
- ◆ [“Congestion Control” on page 203](#)
- ◆ [“Class of Service” on page 213](#)
- ◆ [“Quality of Service” on page 227](#)
- ◆ [“VoIP Traffic Configuration” on page 243](#)
- ◆ [“Security Measures” on page 249](#)
- ◆ [“Basic Administration Protocols” on page 355](#)
- ◆ [“IP Configuration” on page 427](#)
- ◆ [“IP Services” on page 453](#)
- ◆ [“Multicast Filtering” on page 467](#)

3

Basic Management Tasks

This chapter describes the following topics:

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

Displaying System Information

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

Parameters

These parameters are displayed:

- ◆ **System Description** – Brief description of device type.
- ◆ **System Object ID** – MIB II object ID for switch's network management subsystem.
(ECS4210-12P: 1.3.6.1.4.1.259.10.1.42.104,
ECS4210-12T: 1.3.6.1.4.1.259.10.1.42.103,
ECS4210-28T: 1.3.6.1.4.1.259.10.1.42.101,
ECS4210-28P: 1.3.6.1.4.1.259.10.1.42.102)
- ◆ **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.
- ◆ **System EEE** – Enables or disables Energy Efficient Ethernet. When supported by devices on both ends of a link, each side of the link can disable portions of system functionality and save power during periods of low link utilization. Support for EEE is advertised during link-up auto-negotiation.

Web Interface

To configure general system information:

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

Figure 3: System Information

System > General	
System Description	Managed 24G+4GSFP Switch
System Object ID	1.3.6.1.4.1.259.10.1.42.101
System Up Time	0 days, 0 hours, 6 minutes, and 33.53 seconds
System Name	<input type="text"/>
System Location	<input type="text"/>
System Contact	<input type="text"/>
System EEE	<input type="checkbox"/> Enabled

Apply 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.
- ◆ **Internal Power Status** – Displays the status of the internal power supply.

Management Software Information

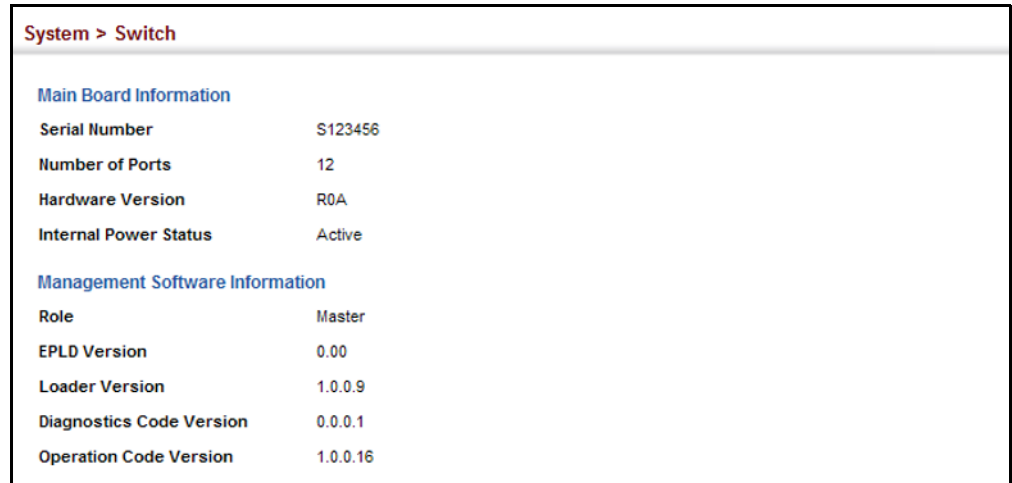
- ◆ **Role** – Shows that this switch is operating as Master or Slave.
- ◆ **EPLD Version** – Version number of Erasable Programmable Logic Device.
- ◆ **Loader Version** – Version number of loader code.
- ◆ **Diagnostics Code Version** – Version of Power-On Self-Test (POST) and boot code.
- ◆ **Operation Code Version** – Version number of runtime code.

Web Interface

To view hardware and software version information.

1. Click System, then Switch.

Figure 4: General Switch Information



System > Switch	
Main Board Information	
Serial Number	S123456
Number of Ports	12
Hardware Version	R0A
Internal Power Status	Active
Management Software Information	
Role	Master
EPLD Version	0.00
Loader Version	1.0.0.9
Diagnostics Code Version	0.0.0.1
Operation Code Version	1.0.0.16

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. Compared to standard Ethernet frames that run only up to 1.5 KB, using jumbo frames significantly reduces the per-packet overhead required to process protocol encapsulation fields.

Usage Guidelines

To use jumbo frames, both the source and destination end nodes (such as a computer or server) must support this feature. Also, when the connection is operating at full duplex, all switches in the network between the two end nodes must be able to accept the extended frame size. And for half-duplex connections, all devices in the collision domain would need to support jumbo frames.

Parameters

The following parameters are displayed:


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

Web Interface

To configure support for jumbo frames:

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

Figure 5: Configuring Support for Jumbo Frames



Displaying Bridge Extension Capabilities

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

Parameters

The following parameters are displayed:

- ◆ **Extended Multicast Filtering Services** – This switch does not support the filtering of individual multicast addresses based on GMRP (GARP Multicast Registration Protocol).
- ◆ **Traffic Classes** – This switch provides mapping of user priorities to multiple traffic classes. (Refer to [“Class of Service” on page 213.](#))
- ◆ **Static Entry Individual Port** – This switch allows static filtering for unicast and multicast addresses. (Refer to [“Setting Static Addresses” on page 165.](#))
- ◆ **VLAN Version Number** – Based on IEEE 802.1Q, “1” indicates Bridges that support only single spanning tree (SST) operation, and “2” indicates Bridges that support multiple spanning tree (MST) operation.
- ◆ **VLAN Learning** – This switch uses Independent VLAN Learning (IVL), where each port maintains its own filtering database.
- ◆ **Local VLAN Capable** – This switch does not support multiple local bridges outside of the scope of 802.1Q defined VLANs.

- ◆ **Configurable PVID Tagging** – This switch allows you to override the default Port VLAN ID (PVID used in frame tags) and egress status (VLAN-Tagged or Untagged) on each port. (Refer to “VLAN Configuration” on page 135.)
- ◆ **Max Supported VLAN Numbers** – The maximum number of VLANs supported on this switch.
- ◆ **Max Supported VLAN ID** – The maximum configurable VLAN identifier supported on this switch.
- ◆ **GMRP** – GARP Multicast Registration Protocol (GMRP) allows network devices to register end stations with multicast groups. This switch does not support GMRP; it uses the Internet Group Management Protocol (IGMP) to provide automatic multicast filtering.

Web Interface

To view Bridge Extension information:

1. Click System, then Capability.

Figure 6: Displaying Bridge Extension Configuration

The screenshot shows the 'System > Capability' configuration page. It is divided into two sections: 'General Capability' and 'Bridge Extension'. Under 'General Capability', 'Jumbo Frame' is set to 'Enabled'. Under 'Bridge Extension', the following settings are displayed:

Setting	Value
Extended Multicast Filtering Services	No
Traffic Classes	Enabled
Static Entry Individual Port	Yes
VLAN Version Number	1
VLAN Learning	IVL
Local VLAN Capable	No
Configurable PVID Tagging	Yes
Max Supported VLAN Numbers	4094
Max Supported VLAN ID	4094
GMRP	Disabled

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

Parameters

The following parameters are displayed:

- ◆ **Copy Type** – The firmware copy operation includes these options:
 - FTP Upgrade – Copies a file from an FTP server to the switch.
 - FTP Download – Copies a file from the switch to an FTP server.
 - HTTP Upgrade – Copies a file from a management station to the switch.
 - HTTP Download – Copies a file from the switch to a management station
 - TFTP Upgrade – Copies a file from a TFTP server to the switch.
 - TFTP Download – Copies a file from the switch to a TFTP server.
- ◆ **FTP/TFTP Server IP Address** – The IP address of an FTP/TFTP server.
- ◆ **User Name** – The user name for FTP server access.
- ◆ **Password** – The password for FTP server access.
- ◆ **File Type** – Specify Operation Code or Loader.
- ◆ **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 128 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 TFTP 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 Upgrade, HTTP Upgrade, or TFTP Upgrade as the file transfer method.
4. If FTP or TFTP Upgrade is used, enter the IP address of the file server.
5. If FTP Upgrade is used, enter the user name and password for your account on the FTP server.
6. Set the file type to Operation Code or Loader.
7. Enter the name of the file to download.
8. Select a file on the switch to overwrite or specify a new file name.
9. Then click Apply.

Figure 7: Copy Firmware

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 Upgrade'. The 'TFTP Server IP Address' text box contains '192.168.0.99'. The 'File Type' dropdown is set to 'Operation Code'. The 'Source File Name' text box is empty. The 'Destination File Name' section has a radio button selected for the first option, which is a dropdown menu showing 'ES4328QV_V1.0.0.5.bix'. There is a second radio button for an empty text box. At the bottom right, there are 'Apply' and 'Revert' buttons.

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

Saving the Running Configuration to a Local File

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

Parameters

The following parameters are displayed:

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



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

Web Interface

To save the running configuration file:

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

Figure 8: Saving the Running Configuration

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

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

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

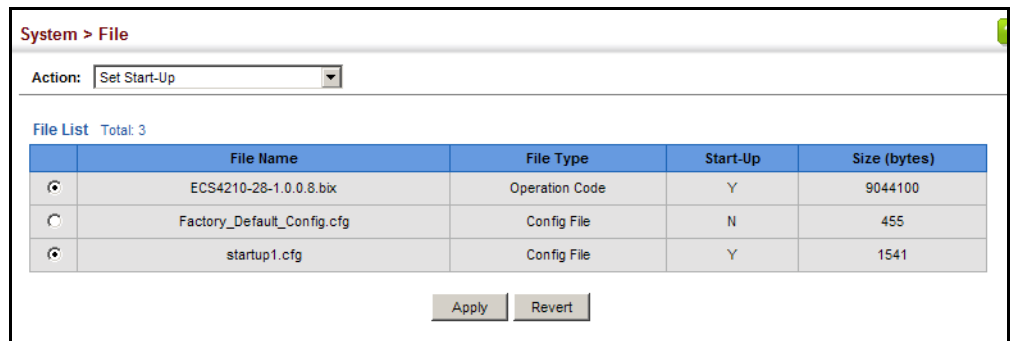
Setting The Start-Up File Use the System > File (Set Start-Up) page to specify the firmware or configuration file to use for system initialization.

Web Interface

To set a file to use for system initialization:

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

Figure 9: Setting Start-Up Files



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.

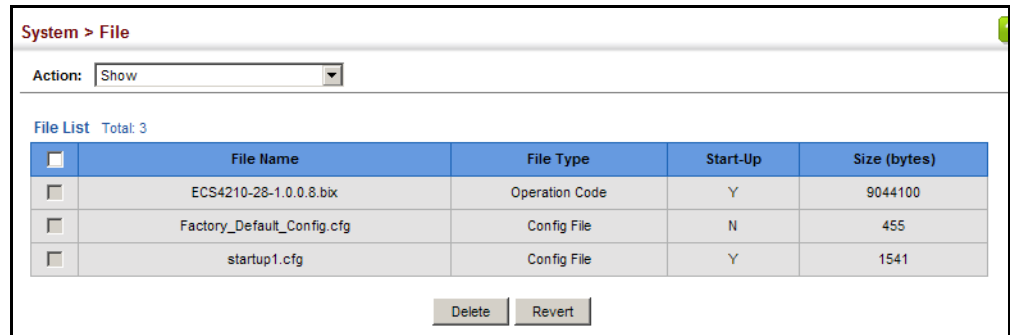
Web Interface

To show the system files:

1. Click System, then File.
2. Select Show from the Action list.

- To delete a file, mark it in the File List and click Delete.

Figure 10: Displaying System Files



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 ECS4210-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 *ECS4210-SERIES.BIX*

from the server even though *ECS4210-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, *ecs4210-series.bix* and *ECS4210-SERIES.bix* are considered to be unique files. Thus, if the upgrade file is stored as *ECS4210-SERIES.bix* on a case-sensitive server, then the switch (requesting *ecs4210-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 *ECS4210-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.
- *password* – Defines the password for the FTP connection. To differentiate the password from the user name and host portions of the URL, a colon (:) must precede the password, and an “at” symbol (@), must follow the password. If the password is omitted, then "" (an empty string) is the assumed password for the connection.
- *host* – Defines the IP address of the FTP server. Valid IP addresses consist of four numbers, 0 to 255, separated by periods. DNS host names are not recognized.
- *filedir* – Defines the directory, relative to the FTP server root, where the upgrade file can be found. Nested directory structures are accepted. The directory name must be separated from the host, and in nested directory structures, from the parent directory, with a prepended forward slash “/”.
- **/** – The forward slash must be the last character of the URL.

Examples

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

- **tftp://192.168.0.1/**
The image file is in the TFTP root directory.
- **tftp://192.168.0.1/switch-opcode/**
The image file is in the “switch-opcode” directory, relative to the TFTP root.

- `tftp://192.168.0.1/switches/opcode/`

The image file is in the “opcode” directory, which is within the “switches” parent directory, relative to the TFTP root.

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

- `ftp://192.168.0.1/`

The user name and password are empty, so “anonymous” will be the user name and the password will be blank. The image file is in the FTP root directory.

- `ftp://switches:upgrade@192.168.0.1/`

The user name is “switches” and the password is “upgrade”. The image file is in the FTP root.

- `ftp://switches:upgrade@192.168.0.1/switches/opcode/`

The user name is “switches” and the password is “upgrade”. The image file is in the “opcode” directory, which is within the “switches” parent directory, relative to the FTP root.

Web Interface

To configure automatic code upgrade:

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

Figure 11: Configuring Automatic Code Upgrade

The screenshot shows a web interface for configuring automatic code upgrade. The breadcrumb path is "System > File". The "Action" dropdown menu is set to "Automatic Operation Code Upgrade". Below this, there is a section for "Automatic Opcode Upgrade" with a checkbox labeled "Enabled" that is currently unchecked. A text input field for "Automatic Upgrade Location URL" is present and empty. A note below the input field states: "Note: For automatic upgrades, the operation code file name must be set as ECS4210-Series.bix." At the bottom right of the form, there are two buttons: "Apply" and "Revert".

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

```
.  
. .  
Automatic Upgrade is looking for a new image  
New image detected: current version 1.0.1.5; new version 1.1.2.0  
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 - Manually) 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 Manually from the Maintain Type list.
4. Enter the time and date in the appropriate fields.
5. Click Apply

Figure 12: Manually Setting the System Clock

The screenshot shows a web interface for configuring the system time. At the top, it says "System > Time". Below that is a "Step:" dropdown menu set to "1. Configure General". The "Current Time" is displayed as "2011-9-26 7:53:34". The "Maintain Type" is set to "Manual". There are input fields for "Hours" (7), "Minutes" (53), and "Seconds" (34). Below these are fields for "Month" (9), "Day" (26), and "Year" (2011). 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 Action list.
3. Select SNTP from the Maintain Type list.
4. Modify the polling interval if required.

5. Click Apply

Figure 13: Setting the Polling Interval for SNTP

The screenshot shows a web interface for configuring system time. At the top, it says "System > Time". Below that is a "Step:" dropdown menu set to "1. Configure General". The "Current Time" is displayed as "2009-9-14 15:21:12". The "Maintain Type" is a dropdown menu set to "SNTP". Under the heading "SNTP Configuration", there is a field for "SNTP Polling Interval (16-16384)" with the value "16" entered and the unit "sec". At the bottom right, there are two buttons: "Apply" and "Revert".

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

Parameters

The following parameters are displayed:

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

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

Web Interface

To set the clock maintenance type to NTP:

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

Figure 14: Configuring NTP

The screenshot shows the 'System > Time' configuration page. At the top, there is a breadcrumb 'System > Time' and a 'Step:' dropdown menu set to '1. Configure General'. Below this, the 'Current Time' is displayed as '2011-9-26 7:57:29'. The 'Maintain Type' is set to 'NTP' via a dropdown menu. Under the 'NTP Configuration' section, the 'Authentication Status' is 'Enabled' (indicated by a checked checkbox), and the 'Polling Interval' is '1024 sec'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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

Specifying SNTP Time Servers

Use the System > Time (Configure Time Server) 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. Enter the IP address of up to three time servers.
4. Click Apply.

Figure 15: Specifying SNTP Time Servers

The screenshot shows the 'System > Time' configuration page at the '2. Configure Time Server' step. The 'Action:' dropdown is set to 'Configure SNTP Server'. There are three input fields for SNTP server IP addresses: 'SNTP Server IP Address 1' with the value '10.1.0.19', 'SNTP Server IP Address 2' with the value '137.62.140.80', and 'SNTP Server IP Address 3' with the value '128.250.36.2'. At the bottom right, there are 'Apply' and 'Revert' buttons.

Specifying NTP Time Servers

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

Parameters

The following parameters are displayed:

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

Web Interface

To add an NTP time server to the server list:

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

Figure 16: Adding an NTP Time Server

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

To show the list of configured NTP time servers:

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

Figure 17: Showing the NTP Time Server List

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

Specifying NTP Authentication Keys

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

Parameters

The following parameters are displayed:

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

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

Web Interface

To add an entry to NTP authentication key list:

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

Figure 18: Adding an NTP Authentication Key

To show the list of configured NTP authentication keys:

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

Figure 19: Showing the NTP Authentication Key List

NTP Authentication Key List		Total: 1
<input type="checkbox"/>	Authentication Key	Key Context
<input type="checkbox"/>	3	8J0774Q6699747D10867F12S505J62770084708278G1357878N8475052H3Q89137L8

Setting the Time Zone

Use the System > Time (Configure Time Server) page to set the time zone. SNTP uses Coordinated Universal Time (or UTC, formerly Greenwich Mean Time, or GMT) based on the time at the Earth's prime meridian, zero degrees longitude, which passes through Greenwich, England. To display a time corresponding to your local time, you must indicate the number of hours and minutes your time zone is west (before) or east (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:

- ◆ **Direction:** Configures the time zone to be before (west of) or after (east of) UTC.
- ◆ **Name** – Assigns a name to the time zone. (Range: 1-29 characters)
- ◆ **Hours** – The number of hours before/after UTC. The maximum value before UTC is 12. The maximum value after UTC is 13. (Range: 0-13 hours)
- ◆ **Minutes** – The number of minutes before/after UTC. (Range: 0-59 minutes)

Web Interface

To set your local time zone:

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

Figure 20: Setting the Time Zone

The screenshot shows a web interface for configuring the system time zone. The breadcrumb is 'System > Time'. The current step is '3. Configure Time Zone'. The 'Direction' dropdown is set to 'After UTC'. The 'Name' text box contains 'UTC'. The 'Hours (0-13)' text box contains '0'. The 'Minutes (0-59)' text box contains '0'. A blue note below the text boxes states: 'Note: The maximum value before UTC is 12:00. The maximum value after UTC is 13:00.' At the bottom right, there are two buttons: 'Apply' and 'Revert'.

Configuring the Console Port

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

Parameters

The following parameters are displayed:

- ◆ **Login Timeout** – Sets the interval that the system waits for a user to log into the CLI. If a login attempt is not detected within the timeout interval, the connection is terminated for the session. (Range: 10-300 seconds; Default: 300 seconds)
- ◆ **Exec Timeout** – Sets the interval that the system waits until user input is detected. If user input is not detected within the timeout interval, the current session is terminated. (Range: 60-65535 seconds; Default: 600 seconds)

- ◆ **Password Threshold** – Sets the password intrusion threshold, which limits the number of failed logon attempts. When the logon attempt threshold is reached, the system interface becomes silent for a specified amount of time (set by the Silent Time parameter) before allowing the next logon attempt. (Range: 1-120; Default: 3 attempts)
- ◆ **Silent Time** – Sets the amount of time the management console is inaccessible after the number of unsuccessful logon attempts has been exceeded. (Range: 1-65535 seconds; Default: Disabled)
- ◆ **Data Bits** – Sets the number of data bits per character that are interpreted and generated by the console port. If parity is being generated, specify 7 data bits per character. If no parity is required, specify 8 data bits per character. (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, Auto; Default: 115200)



Note: Due to a hardware limitation, the terminal program connected to the console port must be set to 8 data bits when using Auto baud rate detection.



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

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

Web Interface

To configure parameters for the console port:

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

Figure 21: Console Port Settings

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

Login Timeout (10-300)	<input type="text" value="300"/> sec
Exec Timeout (60-65535)	<input checked="" type="checkbox"/> <input type="text" value="600"/> sec
Password Threshold (1-120)	<input checked="" type="checkbox"/> <input type="text" value="3"/>
Silent Time (1-65535)	<input type="checkbox"/> <input type="text"/> sec
Data Bits	<input type="text" value="8"/>
Stop Bits	<input type="text" value="1"/>
Parity	<input type="text" value="None"/>
Speed	<input type="text" value="115200"/> baud

Buttons: Apply, Revert

Configuring Telnet Settings

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

Parameters

The following parameters are displayed:

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

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



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

Web Interface

To configure parameters for the console port:

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

Figure 22: Telnet Connection Settings

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

Apply 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 23: Displaying CPU Utilization



Displaying Memory Utilization

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

Parameters

The following parameters are displayed:

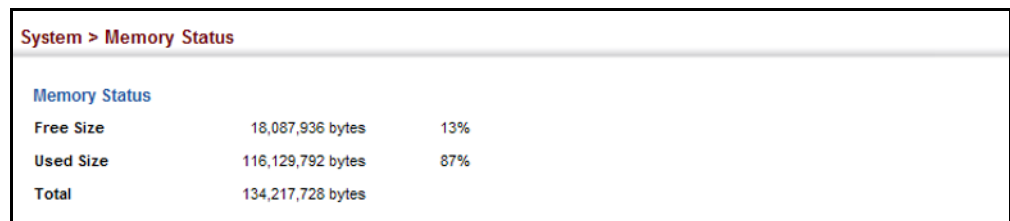
- ◆ **Free Size** – The amount of memory currently free for use.
- ◆ **Used Size** – The amount of memory allocated to active processes.
- ◆ **Total** – The total amount of system memory.

Web Interface

To display memory utilization:

1. Click System, then Memory Status.

Figure 24: Displaying Memory Utilization



The screenshot shows a web interface titled "System > Memory Status". Below the title is a table with the following data:

Memory Status		
Free Size	18,087,936 bytes	13%
Used Size	116,129,792 bytes	87%
Total	134,217,728 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.
- ◆ To retain all configuration information stored in non-volatile memory, click the Save button prior to resetting the 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 as described under [“Saving the Running Configuration to a Local File”](#) on page 73.

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

Figure 25: Restarting the Switch (Immediately)

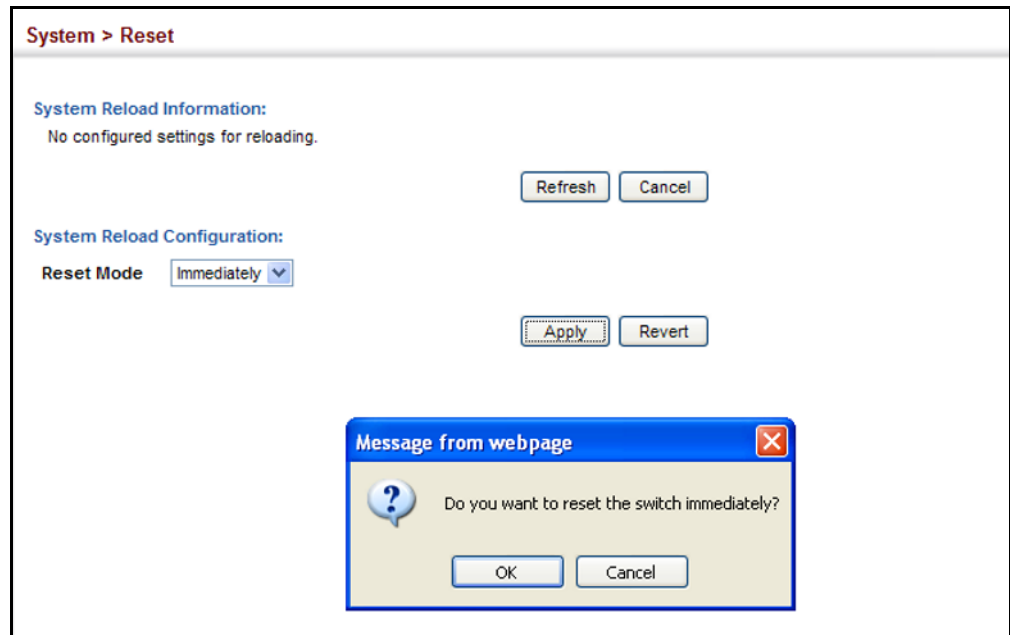


Figure 26: Restarting the Switch (In)

System > Reset

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

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

Figure 27: Restarting the Switch (At)

System > Reset

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

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

Figure 28: Restarting the Switch (Regularly)

System > Reset

System Reload Information:
No configured settings for reloading.

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

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

4

Interface Configuration

This chapter describes the following topics:

- ◆ [Port Configuration](#) – Configures connection settings, including auto-negotiation, or manual setting of speed, duplex mode, and flow control.
- ◆ [Local Port Mirroring](#) – Sets the source and target ports for mirroring on the local switch.
- ◆ [Remote Port Mirroring](#) – Configures mirroring of traffic from remote switches for analysis at a destination port on the local switch.
- ◆ [Displaying Statistics](#) – Shows Interface, Etherlike, and RMON port statistics in table or chart form.
- ◆ [Cable Test](#) – Tests the cable attached to a port.
- ◆ [Trunk Configuration](#) – Configures static or dynamic trunks.
- ◆ [Traffic Segmentation](#) – Configures the uplinks and down links to a segmented group of ports.
- ◆ [VLAN Trunking](#) – Configures a tunnel across one or more intermediate switches which pass traffic for VLAN groups to which they do not belong.

Port Configuration

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

Configuring by Port List

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

Command Usage

- ◆ Auto-negotiation must be disabled before you can configure or force an RJ-45 interface to use the Speed/Duplex mode or Flow Control options.
- ◆ When using auto-negotiation, the optimal settings will be negotiated between the link partners based on their advertised capabilities. To set the speed, duplex

mode, or flow control under auto-negotiation, the required operation modes must be specified in the capabilities list for an interface.

- ◆ The 1000BASE-T standard does not support forced mode. Auto-negotiation should always be used to establish a connection over any 1000BASE-T port or trunk. If not used, the success of the link process cannot be guaranteed when connecting to other types of switches.
- ◆ The Speed/Duplex mode is fixed at 1000full on the Gigabit SFP ports. When auto-negotiation is enabled, the only attributes which can be advertised include flow control and symmetric pause frames.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-12/28)
- ◆ **Type** – Indicates the port type. (1000BASE-T, 1000BASE 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.
- ◆ **Autonegotiation (Port Capabilities)** – Allows auto-negotiation to be enabled/disabled. When auto-negotiation is enabled, you need to specify the capabilities to be advertised. When auto-negotiation is disabled, you can force the settings for speed, mode, and flow control. The following capabilities are supported.
 - **10h** - Supports 10 Mbps half-duplex operation
 - **10f** - Supports 10 Mbps full-duplex operation
 - **100h** - Supports 100 Mbps half-duplex operation
 - **100f** - Supports 100 Mbps full-duplex operation
 - **1000f** - Supports 1000 Mbps full-duplex operation
 - **Sym** (Gigabit only) - Check this item to transmit and receive pause frames.
 - **FC** - Flow control can eliminate frame loss by “blocking” traffic from end stations or segments connected directly to the switch when its buffers fill. When enabled, back pressure is used for half-duplex operation and IEEE 802.3-2005 (formally IEEE 802.3x) for full-duplex operation.

Default: Autonegotiation enabled; Advertised capabilities for
1000BASE-T – 10half, 10full, 100half, 100full, 1000full;
1000Base-SX/LX/LH – 1000full

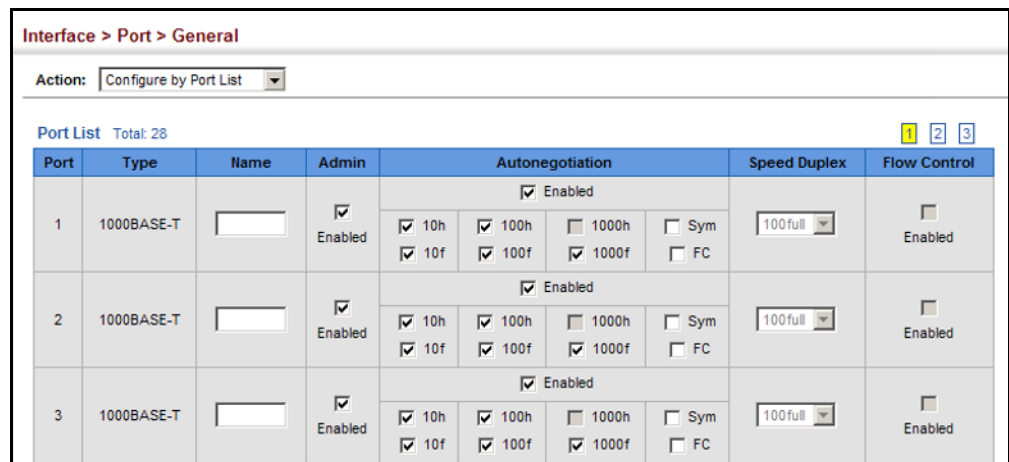
- ◆ **Speed/Duplex** – Allows you to manually set the port speed and duplex mode. (i.e., with auto-negotiation disabled)
- ◆ **Flow Control** – Allows automatic or manual selection of flow control.

Web Interface

To configure port connection parameters:

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

Figure 29: Configuring Connections by Port List



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 95](#).

Web Interface

To configure port connection parameters:

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

5. Click Apply.

Figure 30: Configuring Connections by Port Range

Interface > Port > General

Action:

Port Range (1-28) -

Admin Enabled

Autonegotiation Enabled

10h 100h 1000h Sym

10f 100f 1000f FC

Speed Duplex

Flow Control Enabled

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, 100Base SFP or 1000Base SFP)
- ◆ **Name** – Interface label.
- ◆ **Admin** – Shows if the port is enabled or disabled.
- ◆ **Oper Status** – Indicates if the link is Up or Down.
- ◆ **Media Type** – Not applicable for this switch.
- ◆ **Autonegotiation** – Shows if auto-negotiation is enabled or disabled.
- ◆ **Oper Speed Duplex** – Shows the current speed and duplex mode.
- ◆ **Oper Flow Control** – Shows the flow control type used.

Web Interface

To display port connection parameters:

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

Figure 31: Displaying Port Information

Interface > Port > General

Action: Show Information


Port List Total: 28

Port	Type	Name	Admin	Oper Status	Autonegotiation	Oper Speed Duplex	Oper Flow Control
1	1000BASE-T		Enabled	Up	Enabled	100full	None
2	1000BASE-T		Enabled	Down	Enabled	1000full	None
3	1000BASE-T		Enabled	Down	Enabled	1000full	None
4	1000BASE-T		Enabled	Down	Enabled	1000full	None
5	1000BASE-T		Enabled	Down	Enabled	1000full	None
6	1000BASE-T		Enabled	Down	Enabled	1000full	None
7	1000BASE-T		Enabled	Down	Enabled	1000full	None
8	1000BASE-T		Enabled	Down	Enabled	1000full	None
9	1000BASE-T		Enabled	Down	Enabled	1000full	None
10	1000BASE-T		Enabled	Down	Enabled	1000full	None

Configuring Local Port Mirroring

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

Figure 32: Configuring Local Port Mirroring



Command Usage

- ◆ Traffic can be mirrored from one or more source ports to a destination port on the same switch (local port mirroring as described in this section), 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 101](#)).
- ◆ Monitor port speed should match or exceed source port speed, otherwise traffic may be dropped from the monitor port.
- ◆ When mirroring VLAN traffic (see [“Configuring VLAN Mirroring” on page 162](#)) or packets based on a source MAC address (see [“Configuring MAC Address](#)

Mirroring" on page 173), the target port cannot be set to the same target ports as that used for port mirroring by this command.

- ◆ When traffic matches the rules for both port mirroring, and for mirroring of VLAN traffic or packets based on a MAC address, the matching packets will not be sent to target port specified for port mirroring.
- ◆ Note that Spanning Tree BPDU packets are not mirrored to the target port.
- ◆ The destination port cannot be a trunk or trunk member port.

Parameters

These parameters are displayed:

- ◆ **Source Port** – The port whose traffic will be monitored.
(Range: 1-12/28)
- ◆ **Target Port** – The port that will mirror the traffic on the source port.
(Range: 1-12/28)
- ◆ **Type** – Allows you to select which traffic to mirror to the target port, Rx (receive), Tx (transmit), or Both. (Default: Both)

Web Interface

To configure a local mirror session:

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

Figure 33: Configuring Local Port Mirroring

Interface > Port > Mirror

Action: Add

Source Port Unit 1 Port 7

Target Port Unit 1 Port 8

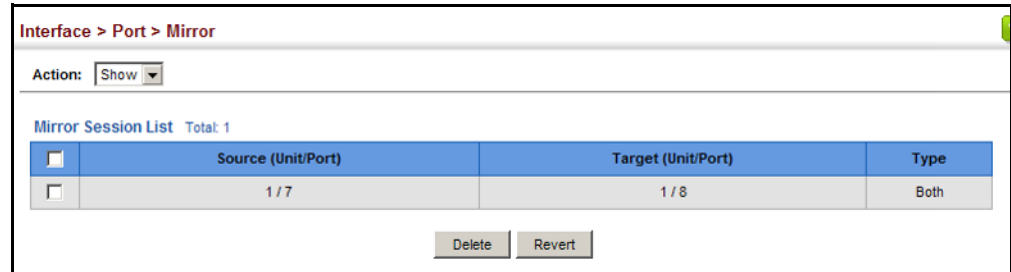
Type Both

Apply Revert

To display the configured mirror sessions:

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

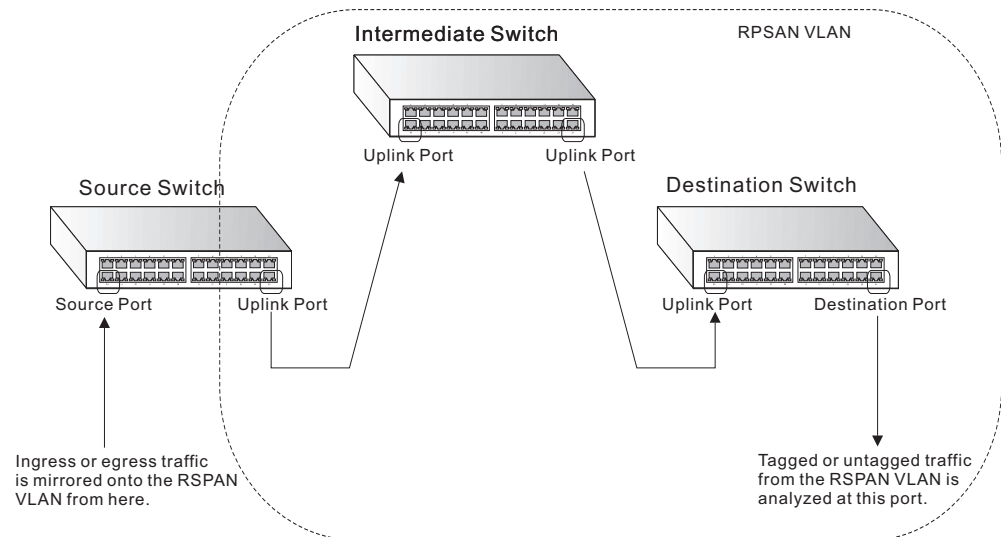
Figure 34: Displaying Local Port Mirror Sessions



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 35: Configuring Remote Port Mirroring



Command Usage

- ◆ Traffic can be mirrored from one or more source ports to a destination port on the same switch (local port mirroring as described in [“Configuring Local Port Mirroring” on page 99](#)), 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 138](#)) to reserve a VLAN for use by RSPAN (marking the “Remote VLAN” field on this page. (Default VLAN 1 is prohibited.)
2. Set up the source switch on the RSPAN configuration page by specifying the mirror session, the switch’s role (Source), the RSPAN VLAN, and the uplink port¹. Then specify the source port(s), and the traffic type to monitor (Rx, Tx or Both).
3. Set up all intermediate switches on the RSPAN configuration page, entering the mirror session, the switch’s role (Intermediate), the RSPAN VLAN, and the uplink port(s).
4. Set up the destination switch on the RSPAN configuration page by specifying the mirror session, the switch’s role (Destination), the destination port¹, whether or not the traffic exiting this port will be tagged or untagged, and the RSPAN VLAN. Then specify each uplink port where the mirrored traffic is being received.

◆ *RSPAN Limitations*

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

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

1. Only 802.1Q trunk or hybrid (i.e., general use) ports can be configured as an RSPAN uplink or destination ports – access ports are not allowed (see [“Adding Static Members to VLANs” on page 140](#)).

RSPAN uplink ports are enabled on the switch, 802.1X cannot be enabled globally.

- *Port Security* – If port security is enabled on any port, that port cannot be set as an RSPAN uplink port, even though it can still be configured as an RSPAN source or destination port. Also, when a port is configured as an RSPAN uplink port, port security cannot be enabled on that port.

Parameters

These parameters are displayed:

- ◆ **Session** – A number identifying this RSPAN session. (Range: 1-2)
Only two mirror sessions are allowed, including both local and remote mirroring. If local mirroring is enabled (see [page 99](#)), then there is only one session available for RSPAN.
- ◆ **Operation Status** – Indicates whether or not RSPAN is currently functioning.
- ◆ **Switch Role** – Specifies the role this switch performs in mirroring traffic.
 - **None** – This switch will not participate in RSPAN.
 - **Source** - Specifies this device as the source of remotely mirrored traffic.
 - **Intermediate** - Specifies this device as an intermediate switch, transparently passing mirrored traffic from one or more sources to one or more destinations.
 - **Destination** - Specifies this device as a switch configured with a destination port which is to receive mirrored traffic for this session.
- ◆ **Remote VLAN** – The VLAN to which traffic mirrored from the source port will be flooded. The VLAN specified in this field must first be reserved for the RSPAN application using the `VLAN > Static` page (see [page 138](#)).
- ◆ **Uplink Port** – A port on any switch participating in RSPAN through which mirrored traffic is passed on to or received from the RSPAN VLAN.
Only one uplink port can be configured on a source switch, but there is no limitation on the number of uplink ports¹ configured on an intermediate or destination switch.

Only destination and uplink ports will be assigned by the switch as members of the RSPAN VLAN. Ports cannot be manually assigned to an RSPAN VLAN through the `VLAN > Static` page. Nor can GVRP dynamically add port members to an RSPAN VLAN. Also, note that the `VLAN > Static (Show)` page will not display any members for an RSPAN VLAN, but will only show configured RSPAN VLAN identifiers.
- ◆ **Type** – Specifies the traffic type to be mirrored remotely. (Options: Rx, Tx, Both)

- ◆ **Destination Port** – Specifies the destination port¹ to monitor the traffic mirrored from the source ports. Only one destination port can be configured on the same switch per session, but a destination port can be configured on more than one switch for the same session. Also note that a destination port can still send and receive switched traffic, and participate in any Layer 2 protocols to which it has been assigned.
- ◆ **Tag** – Specifies whether or not the traffic exiting the destination port to the monitoring device carries the RSPAN VLAN tag.

Web Interface

To configure a remote mirror session:

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

Figure 36: 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 37: 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	<input type="checkbox"/>
2	<input type="checkbox"/>
3	<input type="checkbox"/>
4	<input checked="" type="checkbox"/>
5	<input type="checkbox"/>

Figure 38: 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	<input type="checkbox"/>
2	<input type="checkbox"/>
3	<input type="checkbox"/>
4	<input checked="" type="checkbox"/>
5	<input type="checkbox"/>

Showing Port or Trunk Statistics

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



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

Parameters

These parameters are displayed:

Table 5: 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 5: Port Statistics (Continued)

Parameter	Description
Deferred Transmissions	A count of frames for which the first transmission attempt on a particular interface is delayed because the medium was busy.
Frames Too Long	A count of frames received on a particular interface that exceed the maximum permitted frame size.
Alignment Errors	The number of alignment errors (missynchronized data packets).
FCS Errors	A count of frames received on a particular interface that are an integral number of octets in length but do not pass the FCS check. This count does not include frames received with frame-too-long or frame-too-short error.
SQE Test Errors	A count of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular interface.
Carrier Sense Errors	The number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame.
Internal MAC Receive Errors	A count of frames for which reception on a particular interface fails due to an internal MAC sublayer receive error.
Internal MAC Transmit Errors	A count of frames for which transmission on a particular interface fails due to an internal MAC sublayer transmit error.
<i>RMON Statistics</i>	
Drop Events	The total number of events in which packets were dropped due to lack of resources.
Jabbers	The total number of frames received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS or alignment error.
Fragments	The total number of frames received that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS or alignment error.
Collisions	The best estimate of the total number of collisions on this Ethernet segment.
Received Octets	Total number of octets of data received on the network. This statistic can be used as a reasonable indication of Ethernet utilization.
Received Packets	The total number of packets (bad, broadcast and multicast) received.
Broadcast Packets	The total number of good packets received that were directed to the broadcast address. Note that this does not include multicast packets.
Multicast Packets	The total number of good packets received that were directed to this multicast address.
Undersize Packets	The total number of packets received that were less than 64 octets long (excluding framing bits, but including FCS octets) and were otherwise well formed.
Oversize Packets	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed.
64 Bytes Packets	The total number of packets (including bad packets) received and transmitted that were 64 octets in length (excluding framing bits but including FCS octets).

Table 5: Port Statistics (Continued)

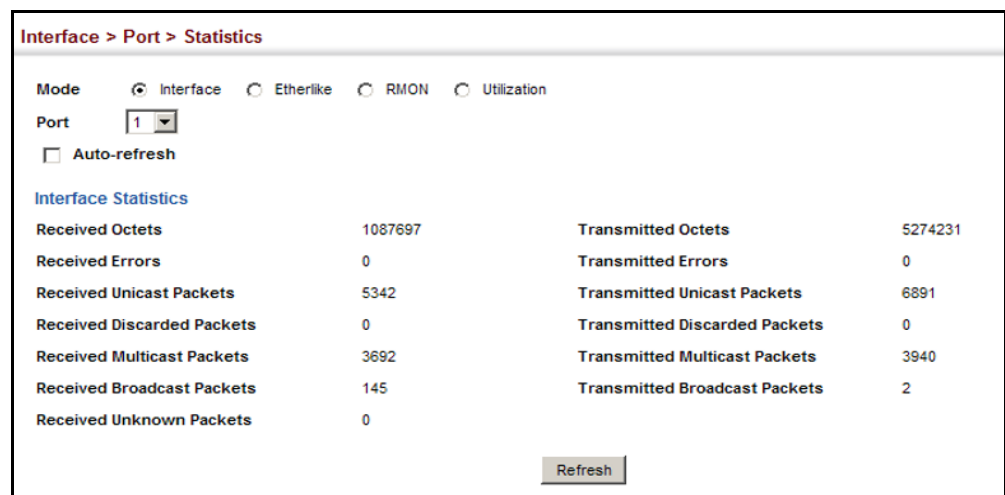
Parameter	Description
65-127 Byte Packets	The total number of packets (including bad packets) received and transmitted where the number of octets fall within the specified range (excluding framing bits but including FCS octets).
128-255 Byte Packets	
256-511 Byte Packets	
512-1023 Byte Packets	
1024-1518 Byte Packets	
1519-1536 Byte Packets	
<i>Utilization Statistics</i>	
Received Octet Rate	Number of octets entering this interface in kbits per second.
Received Packet Rate	Number of packets entering this interface in packets per second.
Received Utilization	The input utilization rate for this interface.
Transmitted Octet Rate	Number of octets leaving this interface in kbits per second.
Transmitted Packet Rate	Number of packets leaving this interface in packets per second.
Transmitted Utilization	The output utilization rate for this interface.

Web Interface

To show a list of port statistics:

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

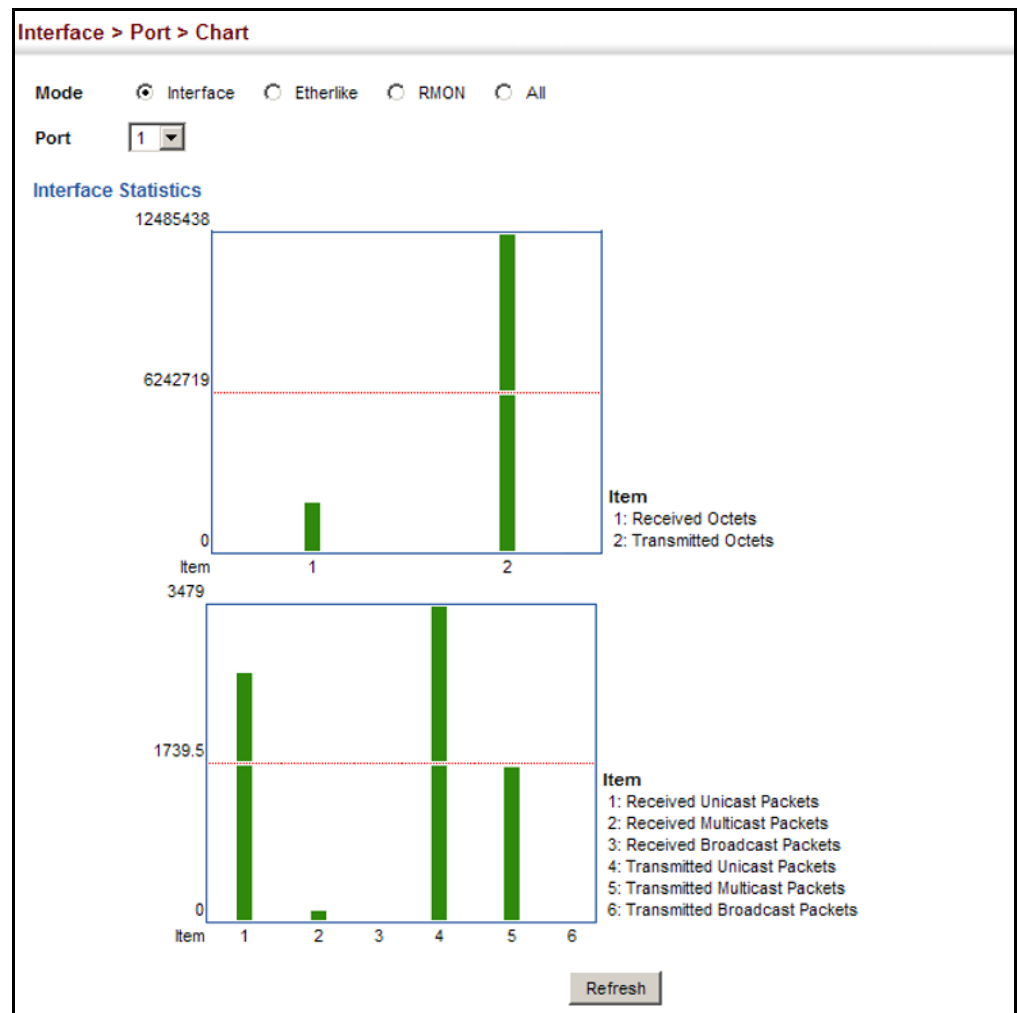
Figure 39: Showing Port Statistics (Table)



To show a chart of port statistics:

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

Figure 40: Showing Port Statistics (Chart)



Performing Cable Diagnostics

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

Command Usage

- ◆ Cable diagnostics are performed using Time Domain Reflectometry (TDR) test methods. TDR analyses the cable by sending a pulsed signal into the cable, and then examining the reflection of that pulse.
- ◆ This cable test is only accurate for Gigabit Ethernet cables 0 - 100 meters long.
- ◆ The test takes approximately 5 seconds. The switch displays the results of the test immediately upon completion, including common cable failures, as well as the status and approximate length to a fault.
- ◆ Potential conditions which may be listed by the diagnostics include:
 - OK: Correctly terminated pair
 - Open: Open pair, no link partner
 - Short: Shorted pair
 - Impedance mismatch: Terminating impedance is not in the reference range.
- ◆ Ports are linked down while running cable diagnostics.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier. (Range: 1-8/24)
- ◆ **Type** – Displays media type. (GE – Gigabit Ethernet)
- ◆ **Link Status** – Shows if the port link is up or down.
- ◆ **Test Result** – The results include common cable failures, as well as the status and approximate distance to a fault, or the approximate cable length if no fault is found.

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

The reported distance to a fault is accurate to +/- 4 meters for most cables. However, depending on cable quality, accuracy may be within +/- 6 meters.

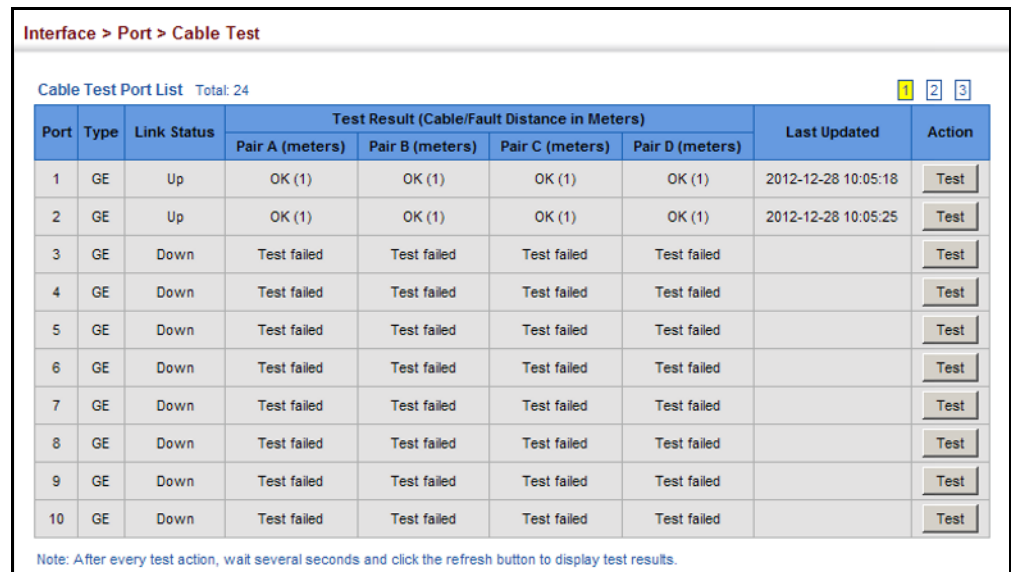
- ◆ **Last Updated** – Shows the last time this port was tested.

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 41: Performing Cable Tests



Interface > Port > Cable Test

Cable Test Port List Total: 24

Port	Type	Link Status	Test Result (Cable/Fault Distance in Meters)				Last Updated	Action
			Pair A (meters)	Pair B (meters)	Pair C (meters)	Pair D (meters)		
1	GE	Up	OK (1)	OK (1)	OK (1)	OK (1)	2012-12-28 10:05:18	Test
2	GE	Up	OK (1)	OK (1)	OK (1)	OK (1)	2012-12-28 10:05:25	Test
3	GE	Down	Test failed	Test failed	Test failed	Test failed		Test
4	GE	Down	Test failed	Test failed	Test failed	Test failed		Test
5	GE	Down	Test failed	Test failed	Test failed	Test failed		Test
6	GE	Down	Test failed	Test failed	Test failed	Test failed		Test
7	GE	Down	Test failed	Test failed	Test failed	Test failed		Test
8	GE	Down	Test failed	Test failed	Test failed	Test failed		Test
9	GE	Down	Test failed	Test failed	Test failed	Test failed		Test
10	GE	Down	Test failed	Test failed	Test failed	Test failed		Test

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

Configuring Port Isolation

Port Isolation can be used to restrict the traffic types or protocol types allowed to pass between specified ports. Isolating traffic to the required uplink and downlink ports can prevent certain types of malicious attacks, and can also reduce the overall amount of traffic crossing the switch.

Enabling Port Isolation Globally

Use the Interface > Port > Isolation (Configure Global) page to enable or disable port isolation globally on the switch.

Parameters

These parameters are displayed:

- ◆ **Port-Isolation Global Status** – Enables or disables port isolation globally on the switch. (Default: Disabled)

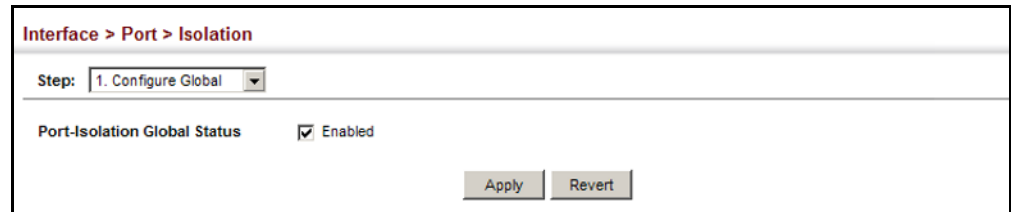
Web Interface

To enable or disable port isolation on the switch:

1. Click Interface, Port, Isolation.
2. Select Configure Global from the Step list.

3. Enable or disable port isolation as required.
4. Click Apply.

Figure 42: Enabling Port Isolation Globally



Interface > Port > Isolation

Step: 1. Configure Global

Port-Isolation Global Status Enabled

Apply Revert

Configuring Port Isolation Profiles

Use the Interface > Port > Isolation (Configure Profile - Add) page to set the traffic type or protocol type to include in a profile.

Command Usage

- ◆ An isolation profile can include any number of traffic or protocol types.
- ◆ An isolation profile must be assigned to an uplink or downlink port using the Configure Interface page before it will take effect.

Parameters

These parameters are displayed:

- ◆ **Profile ID** – Profile identifier. (Range: 1-20)
- ◆ **Traffic Type** – Sets the controlled transmission type. (Options: Unicast, Multicast, Broadcast)
- ◆ **Protocol Type** – Sets the controlled protocol type. (Options: IPv4 ARP, IPv4 DHCP, IPv6 NDP, IPv6 DHCP)

Web Interface

To set the traffic or protocol type to include in a profile:

1. Click Interface, Port, Isolation.
2. Select Configure Profile from the Step list.
3. Select Add from the Action list.
4. Enter a profile identifier, and mark the traffic types or protocol types to control.
5. Click Apply.

Figure 43: Configuring Port Isolation Profiles

To display the configured settings for a profile:

1. Click Interface, Port, Isolation.
2. Select Configure Profile from the Step list.
3. Select Show from the Action list.

Figure 44: Displaying Port Isolation Profiles

Profile ID	Traffic Type	Protocol Type
1		IPv4 DHCP

Assigning Port Isolation Profiles

Use the Interface > Port > Isolation (Configure Interface) page to assign a profile to an uplink or downlink port.

Command Usage

- ◆ When a profile is assigned to a downlink port (i.e., isolated port), traffic matching any of the defined attributes (entering the switch through that port) can only be passed to an uplink port which has been assigned the same profile.
- ◆ Only one profile can be assigned to a downlink port, but one or more profiles can be assigned to an uplink port. This allows a downlink port to pass any traffic defined in its profile to all uplink ports which share the same profile. This also allows an uplink port to service any number of downlink ports which share the same profile.
- ◆ When an uplink port is assigned more than one profile, traffic entering the switch through the uplink port must be checked against all of its assigned profiles to determine the downlink port to which that traffic can be forwarded.

- ◆ When a profile is assigned to a port, any traffic attributes not defined in the profile are subject normal switching rules.

Parameters

These parameters are displayed:

- ◆ **Profile ID** – Profile identifier. (Range: 1-26)
- ◆ **Port** – Port identifier. (Range: 1-12/28)
- ◆ **Uplink Port Group** – Port connected to an upstream interface.
- ◆ **Isolated Port Group** – Port connected to an down stream interface.

Web Interface

To assign a profile to an uplink or downlink port:

1. Click Interface, Port, Isolation.
2. Select Configure Interface from the Step List.
3. Select a profile identifier, and mark the uplink ports or downlink ports to which that profile will the assigned.
4. Click Apply.

Figure 45: Assigning Port Isolation Profiles

Interface > Port > Isolation

Step: 3. Configure Interface

Profile ID: 1

Port List Total: 28

Port	Uplink Port Group	Isolated Port Group
1	<input checked="" type="checkbox"/> Join	<input type="checkbox"/> Join
2	<input type="checkbox"/> Join	<input checked="" type="checkbox"/> Join
3	<input type="checkbox"/> Join	<input checked="" type="checkbox"/> Join
4	<input type="checkbox"/> Join	<input checked="" type="checkbox"/> Join
5	<input type="checkbox"/> Join	<input checked="" type="checkbox"/> Join
6	<input type="checkbox"/> Join	<input type="checkbox"/> Join
7	<input type="checkbox"/> Join	<input type="checkbox"/> Join
8	<input type="checkbox"/> Join	<input type="checkbox"/> Join
9	<input type="checkbox"/> Join	<input type="checkbox"/> Join
10	<input type="checkbox"/> Join	<input type="checkbox"/> Join

Apply Revert

Trunk Configuration

This section describes how to configure static and dynamic trunks.

You can create multiple links between devices that work as one virtual, aggregate link. A 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 8/12 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 8/12 trunks on a switch, with up to eight ports per trunk.
- ◆ The ports at both ends of a connection must be configured as trunk ports.
- ◆ When configuring static trunks on switches of different types, they must be compatible with the Cisco EtherChannel standard.
- ◆ The ports at both ends of a trunk must be configured in an identical manner, including communication mode (i.e., speed, duplex mode and flow control), VLAN assignments, and CoS settings.
- ◆ Any of the Gigabit ports on the front panel can be trunked together, including ports of different media types.
- ◆ All the ports in a trunk have to be treated as a whole when moved from/to, added or deleted from a VLAN.
- ◆ STP, VLAN, and IGMP settings can only be made for the entire trunk.

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

Figure 46: 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: 1-8/12)
- ◆ **Member** – The initial trunk member. Use the Add Member page to configure additional members.
 - **Unit** – Unit identifier. (Range: 1)
 - **Port** – Port identifier. (Range: 1-12/28)

Web Interface

To create a static trunk:

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

Figure 47: Creating Static Trunks

Interface > Trunk > Static

Step: 1. Configure Trunk Action: Add

Trunk ID (1-8) 1

Member Unit 1 Port 5

Apply Revert

To add member ports to a static trunk:

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

Figure 48: Adding Static Trunks Members

Interface > Trunk > Static

Step: 1. Configure Trunk Action: Add Member

Trunk 1

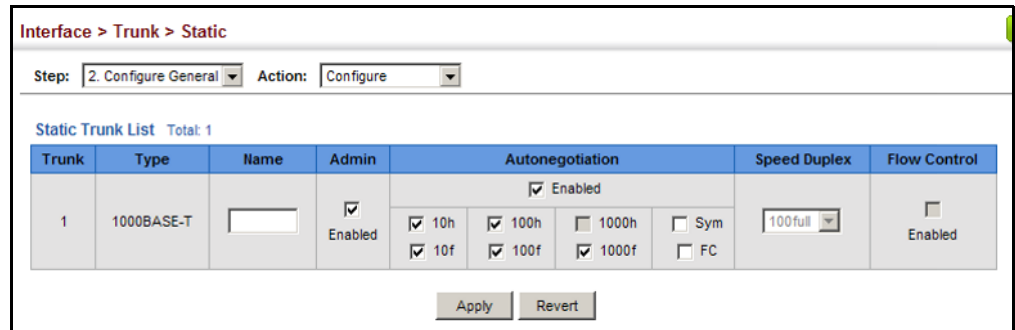
Member Unit 1 Port 1

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


Figure 49: Configuring Connection Parameters for a Static Trunk



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 50: Showing Information for Static Trunks



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 51: 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)

Configure Aggregation Port - General

- ◆ **Port** – Port identifier. (Range: 1-12/28)
- ◆ **LACP Status** – Enables or disables LACP on a port.

Configure Aggregation Port - Actor/Partner

- ◆ **Port** – Port number. (Range: 1-12/28)
- ◆ **Admin Key** – The LACP administration key must be set to the same value for ports that belong to the same LAG. (Range: 0-65535; Default – Actor: 1, Partner: 0)

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

- ◆ **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 for the required LACP group.
4. Click Apply.

Figure 52: Configuring the LACP Aggregator Admin Key

Interface > Trunk > Dynamic

Step: 1. Configure Aggregator

Trunk List Total: 8

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

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 53: 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	<input checked="" type="checkbox"/> Enabled
2	<input type="checkbox"/> Enabled
3	<input type="checkbox"/> Enabled
4	<input type="checkbox"/> Enabled
5	<input type="checkbox"/> 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 54: 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 (0-85535)	System Priority (0-65535)	Port Priority (0-85535)
1	1	32768	32768
2	1	32768	32768
3	1	32768	32768
4	1	32768	32768
5	1	32768	32768

To configure the 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. (Refer to [“Configuring by Port List” on page 95](#) for a description of the parameters.)
5. Click Apply.

Figure 55: Configuring Connection Parameters 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
2	1000BASE-T		<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> Enabled				100full	<input checked="" type="checkbox"/> Enabled
				<input checked="" type="checkbox"/> 10h	<input checked="" type="checkbox"/> 100h	<input type="checkbox"/> 1000h	<input type="checkbox"/> Sym		
				<input checked="" type="checkbox"/> 10f	<input checked="" type="checkbox"/> 100f	<input checked="" type="checkbox"/> 1000f	<input type="checkbox"/> FC		

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



To show the port members of dynamic trunks:

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

Figure 57: Showing Members of Dynamic Trunks



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 6: LACP Port Counters

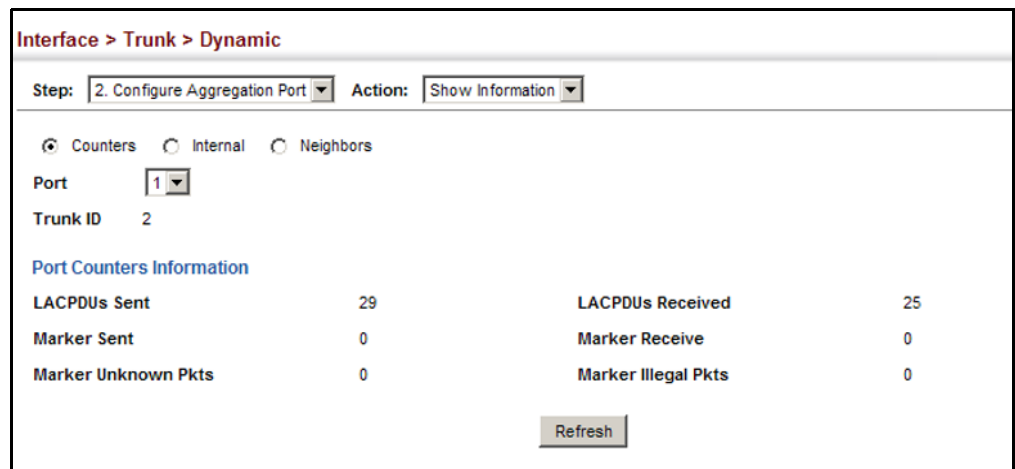
Parameter	Description
LACPDUs Sent	Number of valid LACPDUs transmitted from this channel group.
LACPDUs Received	Number of valid LACPDUs received on this channel group.
Marker Sent	Number of valid Marker PDUs transmitted from this channel group.
Marker Received	Number of valid Marker PDUs received by 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 58: Displaying LACP Port Counters



Displaying LACP Settings and Status for the Local Side

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

Parameters

These parameters are displayed:

Table 7: LACP Internal Configuration Information

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

Table 7: LACP Internal Configuration Information (Continued)

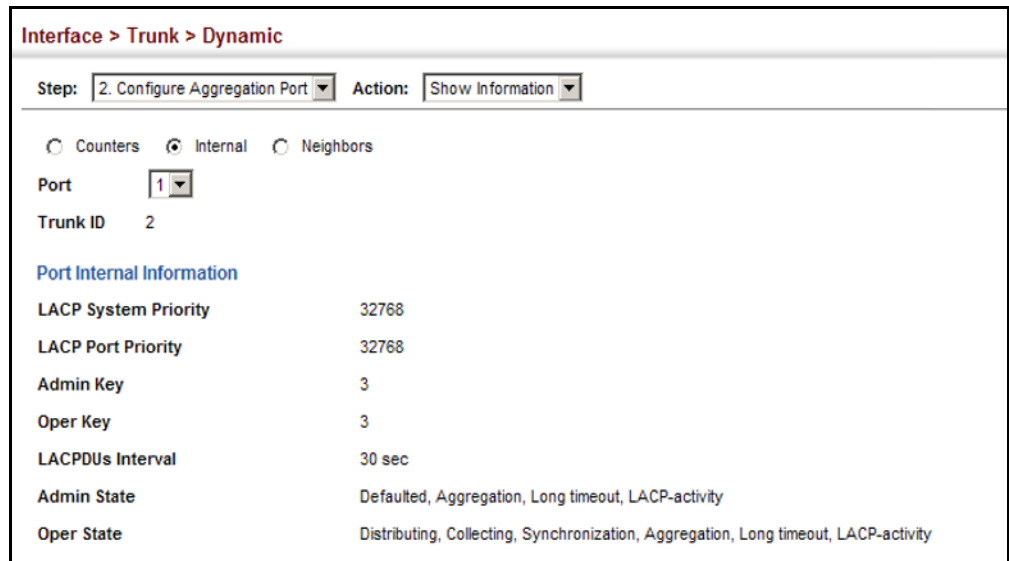
Parameter	Description
Admin State, Oper State	<p>Administrative or operational values of the actor's state parameters:</p> <ul style="list-style-type: none"> ◆ Expired – The actor's receive machine is in the expired state; ◆ Defaulted – The actor's receive machine is using defaulted operational partner information, administratively configured for the partner. ◆ Distributing – If false, distribution of outgoing frames on this link is disabled; i.e., distribution is currently disabled and is not expected to be enabled in the absence of administrative changes or changes in received protocol information. ◆ Collecting – Collection of incoming frames on this link is enabled; i.e., collection is currently enabled and is not expected to be disabled in the absence of administrative changes or changes in received protocol information. ◆ Synchronization – The System considers this link to be IN_SYNC; i.e., it has been allocated to the correct Link Aggregation Group, the group has been associated with a compatible Aggregator, and the identity of the Link Aggregation Group is consistent with the System ID and operational Key information transmitted. ◆ Aggregation – The system considers this link to be aggregatable; i.e., a potential candidate for aggregation. ◆ Long timeout – Periodic transmission of LACPDUs uses a slow transmission rate. ◆ LACP-Activity – Activity control value with regard to this link. (0: Passive; 1: Active)

Web Interface

To display LACP settings and status for the local side:

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

Figure 59: Displaying LACP Port Internal Information



Displaying LACP Settings and Status for the Remote Side

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

Parameters

These parameters are displayed:

Table 8: LACP Remote Device Configuration Information

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

Web Interface

To display LACP settings and status for the remote side:

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

Figure 60: 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

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 61: Enabling Traffic Segmentation

Configuring Uplink and Downlink Ports

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

Command Usage

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

Table 9: 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 128](#)).

- ◆ 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: 1-12/28)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-8/12)
- ◆ **Direction** – Adds an interface to the segmented group by setting the direction to uplink or downlink. (Default: None)

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 62: Configuring Members for Traffic Segmentation

The screenshot shows a web interface titled "Interface > Traffic Segmentation". At the top, there is a "Step:" dropdown menu set to "2. Configure Session" and an "Action:" dropdown menu set to "Add". Below this, there are several input fields and dropdown menus:


- Session ID (1-4)**: A text input field.
- Direction**: A dropdown menu currently set to "Uplink".
- Interface**: Two radio button options:
 - Port (1-28) [] - []
 - Trunk (1-12) [] - []

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 63: Showing Traffic Segmentation Members



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 64: 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 140](#)). 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: 1-12/28)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-8/12)
- ◆ **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 65: Configuring VLAN Trunking

Interface > VLAN Trunking ?

Interface Port Trunk

Port VLAN Trunking List Total: 28 1 2 3

Port	VLAN Trunking Status
1	<input checked="" type="checkbox"/> Enabled
2	<input type="checkbox"/> Enabled
3	<input type="checkbox"/> Enabled
4	<input type="checkbox"/> Enabled
5	<input type="checkbox"/> Enabled
6	<input type="checkbox"/> Enabled
7	<input type="checkbox"/> Enabled
8	<input type="checkbox"/> Enabled
9	<input type="checkbox"/> Enabled
10	<input type="checkbox"/> Enabled

VLAN Configuration

This chapter includes the following topics:

- ◆ [IEEE 802.1Q VLANs](#) – Configures static and dynamic VLANs.
- ◆ [IEEE 802.1Q Tunneling](#) – Configures QinQ tunneling to maintain customer-specific VLAN and Layer 2 protocol configurations across a service provider network, even when different customers use the same internal VLAN IDs.
- ◆ [Protocol VLANs](#) – Configures VLAN groups based on specified protocols.
- ◆ [IP Subnet VLANs](#) – 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](#) – Maps untagged ingress frames to a specified VLAN if the source MAC address is found in the IP MAC address-to-VLAN mapping table.
- ◆ [VLAN Mirroring](#) – Mirrors traffic from one or more source VLANs to a target port.

IEEE 802.1Q VLANs

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

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

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

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

This switch supports the following VLAN features:

- ◆ Up to 256 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 66: 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 140](#)). But you can still enable GVRP on these edge switches, as well as on the core switches in the network.

Figure 67: 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 101](#)). To propagate information about VLAN groups used on this switch to external network devices, you must specify a VLAN ID for each of these groups.

Parameters

These parameters are displayed:

Add

- ◆ **VLAN ID** – ID of VLAN or range of VLANs (1-4094).

Up to 256 VLAN groups can be defined. 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 101](#)).

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.

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 101](#)).

Web Interface

To create VLAN groups:

1. Click VLAN, Static.
2. Select Add from the Action list.
3. Enter a VLAN ID or range of IDs.
4. Mark Enabled to configure the VLAN as operational.
5. Specify whether the VLANs are to be used for remote port mirroring.
6. Click Apply.

Figure 68: Creating Static VLANs

VLAN > Static

Action: Add

VLAN ID (1-4094) 2 -

Status Enabled

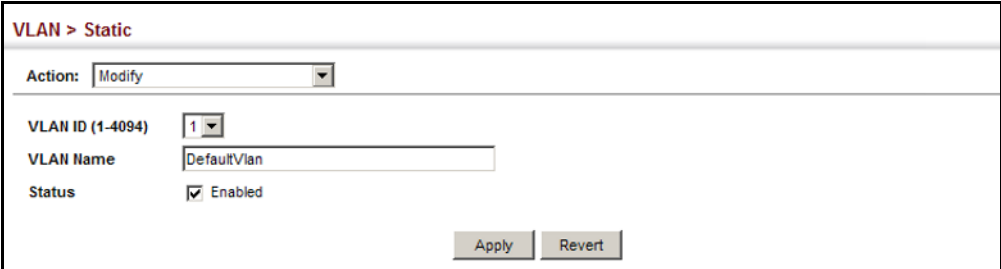
Remote VLAN Enabled

Apply 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. Click Apply.

Figure 69: Modifying Settings for Static VLANs



VLAN > Static

Action:

VLAN ID (1-4094)

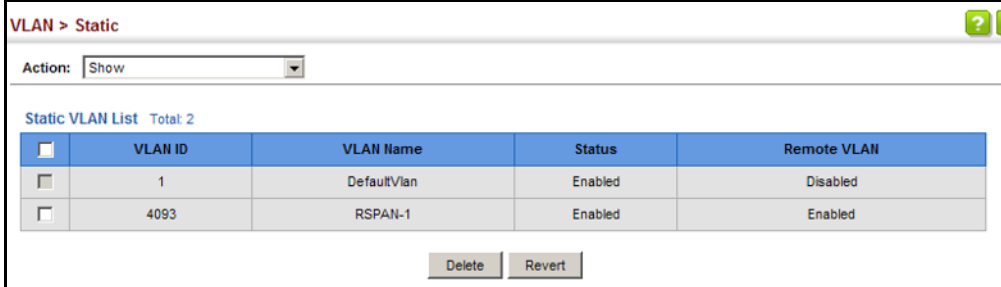
VLAN Name

Status Enabled

To show the configuration settings for VLAN groups:

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

Figure 70: Showing Static VLANs



VLAN > Static

Action:

Static VLAN List Total: 2

<input type="checkbox"/>	VLAN ID	VLAN Name	Status	Remote VLAN
<input type="checkbox"/>	1	DefaultVlan	Enabled	Disabled
<input type="checkbox"/>	4093	RSPAN-1	Enabled	Enabled

Adding Static Members to VLANs

Use the VLAN > Static (Modify VLAN and Member Ports, 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: 1-12/28)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-8/12)
- ◆ **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 131](#)). 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 137.
 - **None:** Interface is not a member of the VLAN. Packets associated with this VLAN will not be transmitted by the interface.



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

Edit Member by Interface

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

Edit Member by Interface Range

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

- ◆ **Port Range** – Displays a list of ports. (Range: 1-12/28)
- ◆ **Trunk Range** – Displays a list of ports. (Range: 1-8/12)



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

Web Interface

To configure static members by the VLAN index:

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

Figure 71: Configuring Static Members by VLAN Index

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

To configure static members by interface:

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

Figure 72: 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	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
2	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
3	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
4	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>

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 73: Configuring Static VLAN Members by Interface Range

VLAN > Static

Action: Edit Member by Interface Range

Interface: Port Trunk

Port Range (1-28): 9 - 10

Mode: Hybrid

VLAN ID (1-4093): 2 -

Membership Type: Tagged Untagged Forbidden None

Apply Revert

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

Parameters

These parameters are displayed:

Configure General

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

Configure Interface

- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Port** – Port Identifier. (Range: 1-12/28)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-8/12)
- ◆ **GVRP Status** – Enables/disables GVRP for the interface. GVRP must be globally enabled for the switch before this setting can take effect (using the Configure General page). When disabled, any GVRP packets received on this port will be discarded and no GVRP registrations will be propagated from other ports. (Default: Disabled)

GVRP cannot be enabled for ports set to Access mode (see [“Adding Static Members to VLANs”](#) on page 140).

- ◆ **GVRP Timers** – Timer settings must follow this rule:
 $3 \times (\text{join timer}) < \text{leave timer} < \text{leaveAll timer}$
 - **Join** – The interval between transmitting requests/queries to participate in a VLAN group. (Range: 20-1000 centiseconds; Default: 20)
 - **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)
 - **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)

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 74: Configuring Global Status of GVRP



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 75: 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	<input type="checkbox"/> Enabled	20	60	1000
2	<input type="checkbox"/> Enabled	20	60	1000
3	<input checked="" type="checkbox"/> Enabled	20	60	1000
4	<input type="checkbox"/> Enabled	20	60	1000
5	<input type="checkbox"/> 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 76: Showing Dynamic VLANs Registered on the Switch

VLAN > Dynamic

Step: 3. Show Dynamic VLAN Action: Show VLAN

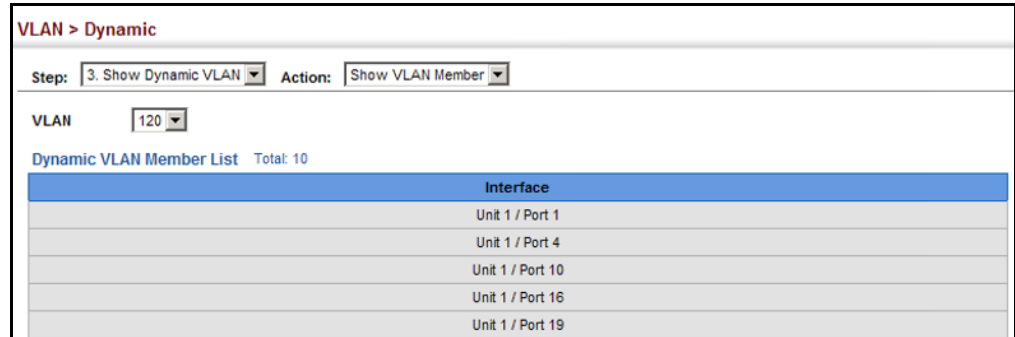
Dynamic VLAN List Total: 2

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 77: Showing the Members of a Dynamic VLAN



The screenshot shows a network configuration page titled "VLAN > Dynamic". At the top, there are two dropdown menus: "Step: 3. Show Dynamic VLAN" and "Action: Show VLAN Member". Below this, a "VLAN" dropdown menu is set to "120". Underneath, it says "Dynamic VLAN Member List Total: 10". A table with a blue header "Interface" lists five members: "Unit 1 / Port 1", "Unit 1 / Port 4", "Unit 1 / Port 10", "Unit 1 / Port 16", and "Unit 1 / Port 19".

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.

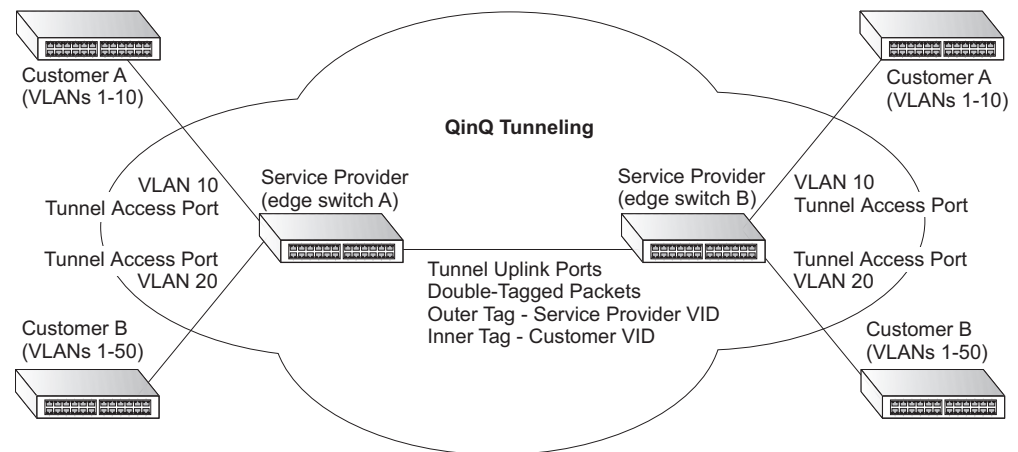
QinQ tunneling uses a single Service Provider VLAN (SPVLAN) for customers who have multiple VLANs. Customer VLAN IDs are preserved and traffic from different customers is segregated within the service provider's network even when they use the same customer-specific VLAN IDs. QinQ tunneling expands VLAN 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 78: 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). This outer tag is used for learning and switching packets across the metropolitan network. 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 152.](#))
2. Create a Service Provider VLAN, also referred to as an SPVLAN (see [“Configuring VLAN Groups” on page 138.](#))
3. Configure the QinQ tunnel access port to Access mode (see [“Adding an Interface to a QinQ Tunnel” on page 153.](#))
4. Configure the QinQ tunnel access port to join the SPVLAN as an untagged member (see [“Adding Static Members to VLANs” on page 140.](#))

5. Configure the SPVLAN ID as the native VID on the QinQ tunnel access port (see [“Adding Static Members to VLANs” on page 140](#)).
6. Configure the QinQ tunnel uplink port to Uplink mode (see [“Adding an Interface to a QinQ Tunnel” on page 153](#)).
7. Configure the QinQ tunnel uplink port to join the SPVLAN as a tagged member (see [“Adding Static Members to VLANs” on page 140](#)).

Enabling QinQ Tunneling on the Switch

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

Parameters

These parameters are displayed:

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

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

The specified ethertype only applies to ports configured in Uplink mode (see [“Adding an Interface to a QinQ Tunnel” on page 153](#)). 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.

All ports on the switch will be set to the same ethertype.

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 79: Enabling QinQ Tunneling

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

Command Usage

- ◆ Use the Configure Global page to set the switch to QinQ mode before configuring a tunnel access port or tunnel uplink port (see [“Enabling QinQ Tunneling on the Switch” on page 152](#)). 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:

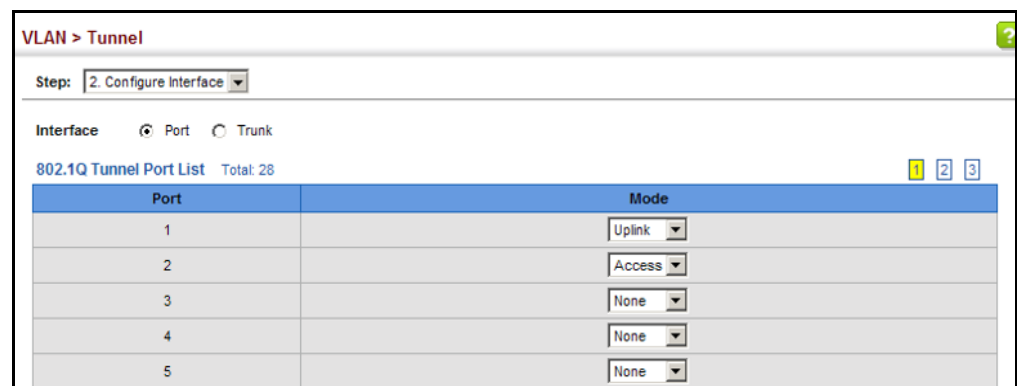
- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Port** – Port Identifier. (Range: 1-12/28)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-8/12)
- ◆ **Mode** – Sets the VLAN membership mode of the port.
 - **None** – The port operates in its normal VLAN mode. (This is the default.)
 - **Access** – 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 80: Adding an Interface to a QinQ Tunnel



Protocol VLANs

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

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

Command Usage

◆ To configure protocol-based VLANs, follow these steps:

1. First configure VLAN groups for the protocols you want to use (see [“Configuring VLAN Groups” on page 138](#)). 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, and protocol-based VLANs are supported concurrently, priority is applied in this sequence, and then port-based VLANs last.

Configuring Protocol VLAN Groups

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

Parameters

These parameters are displayed:

- ◆ **Frame Type** – Choose either Ethernet, RFC 1042, or LLC Other as the frame type used by this protocol.
- ◆ **Protocol Type** – Specifies the protocol type to match. The available options are IP, ARP, RARP, PPPoE Discovery, PPPoE Session, 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 81: 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 82: Displaying Protocol VLANs

VLAN > Protocol

Step: 1. Configure Protocol Action: Show

Protocol to Group Mapping Table Total: 5

<input type="checkbox"/>	Frame Type	Protocol Type	Protocol Group ID
<input type="checkbox"/>	Ethernet	08 06	1
<input type="checkbox"/>	Ethernet	80 35	2
<input type="checkbox"/>	RFC 1042	08 00	1
<input type="checkbox"/>	RFC 1042	80 35	3
<input type="checkbox"/>	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 140), 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: 1-12/28)
- ◆ **Trunk** – Trunk Identifier. (Range: 1-8/12)
- ◆ **Protocol Group ID** – Protocol Group ID assigned to the Protocol VLAN Group. (Range: 1-2147483647)
- ◆ **VLAN ID** – VLAN to which matching protocol traffic is forwarded. (Range: 1-4094)
- ◆ **Priority** – The priority assigned to untagged ingress traffic. (Range: 0-7, where 7 is the highest priority)

Web Interface

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

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

Figure 83: Assigning Interfaces to Protocol VLANs

VLAN > Protocol

Step: 2. Configure Interface Action: Add

Interface Port 3 Trunk

Protocol Group ID 1

VLAN ID (1-4094) 2

Priority (0-7) 0

Apply Revert

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

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

Figure 84: Showing the Interface to Protocol Group Mapping

VLAN > Protocol

Step: 2. Configure Interface Action: Show

Interface Port 2 Trunk 1

Port To Protocol Group Mapping Table Total: 2

<input type="checkbox"/>	Protocol Group ID	VLAN ID	Priority
<input type="checkbox"/>	1	10	0
<input type="checkbox"/>	3	20	0

Delete Revert

Configuring IP Subnet VLANs

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

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

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

Command Usage

- ◆ Each IP subnet can be mapped to only one VLAN ID. An IP subnet consists of an IP address and a mask. The specified VLAN need not be an existing VLAN.
- ◆ When an untagged frame is received by a port, the source IP address is checked against the IP subnet-to-VLAN mapping table, and if an entry is found, the corresponding VLAN ID is assigned to the frame. If no mapping is found, the PVID of the receiving port is assigned to the frame.
- ◆ The IP subnet cannot be a broadcast or multicast IP address.
- ◆ When MAC-based, IP subnet-based, and 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 85: Configuring IP Subnet VLANs

VLAN > IP Subnet

Action: Add

IP Address: 192.168.1.19

Subnet Mask: 255.255.255.0

VLAN (1-4094): 10

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 86: Showing IP Subnet VLANs

VLAN > IP Subnet

Action: Show

IP Subnet to VLAN Mapping Table Total: 1

<input type="checkbox"/>	IP Address	Subnet Mask	VLAN	Priority
<input type="checkbox"/>	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, and 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. Configured MAC addresses can only be unicast addresses. 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)
- ◆ **Mask** – Identifies a range of MAC addresses. The mask can be specified in the format xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx, where an equivalent binary value “1” means relevant and “0” means ignore.

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.
4. Enter an identifier in the VLAN field. Note that the specified VLAN need not already be configured.
5. Click Apply.

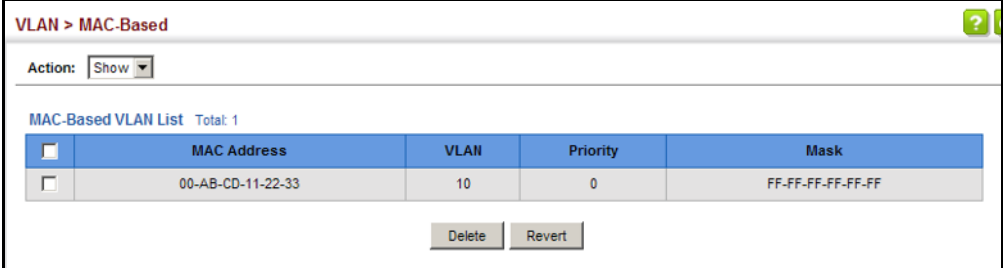
Figure 87: Configuring MAC-Based VLANs

The screenshot shows a web interface titled "VLAN > MAC-Based". At the top, there is an "Action:" dropdown menu with "Add" selected. Below this are four input fields: "MAC Address", "VLAN (1-4094)", "Priority (0-7)", and "Mask". The "Mask" field is pre-filled with the hexadecimal value "ff-ff-ff-ff-ff-ff". At the bottom right of the form, there 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 88: Showing MAC-Based VLANs



The screenshot shows a web interface for configuring VLANs. At the top, it says "VLAN > MAC-Based". Below that, there is an "Action:" dropdown menu set to "Show". Underneath, it says "MAC-Based VLAN List Total: 1". A table with the following columns is displayed: "MAC Address", "VLAN", "Priority", and "Mask". There is a checkbox in the first column of the table. The table contains one row with the following data: MAC Address: 00-AB-CD-11-22-33, VLAN: 10, Priority: 0, Mask: FF-FF-FF-FF-FF-FF. Below the table are "Delete" and "Revert" buttons.

<input type="checkbox"/>	MAC Address	VLAN	Priority	Mask
<input type="checkbox"/>	00-AB-CD-11-22-33	10	0	FF-FF-FF-FF-FF-FF

Configuring VLAN Mirroring

Use the VLAN > Mirror (Add) page to mirror traffic from one or more source VLANs to a target port for real-time analysis. You can then attach a logic analyzer or RMON probe to the target port and study the traffic crossing the source VLAN(s) in a completely unobtrusive manner.

Command Usage

- ◆ All active ports in a source VLAN are monitored for ingress traffic only.
- ◆ All VLAN mirror sessions must share the same target port, preferably one that is not a member of the source VLAN.
- ◆ When VLAN mirroring and port mirroring are both enabled, they must use the same target port.
- ◆ When VLAN mirroring and port mirroring are both enabled, the target port can receive a mirrored packet twice; once from the source mirror port and again from the source mirrored VLAN.
- ◆ The target port receives traffic from all monitored source VLANs and can become congested. Some mirror traffic may therefore be dropped from the target port.
- ◆ When mirroring VLAN traffic or packets based on a source MAC address (see [“Configuring MAC Address Mirroring” on page 173](#)), the target port cannot be set to the same target ports as that used for port mirroring (see [“Configuring Local Port Mirroring” on page 99](#)).
- ◆ When traffic matches the rules for both port mirroring, and for mirroring of VLAN traffic or packets based on a MAC address, the matching packets will not be sent to target port specified for port mirroring.

Parameters

These parameters are displayed:

- ◆ **Source VLAN** – A VLAN whose traffic will be monitored. (Range: 1-4094)
- ◆ **Target Port** – The destination port that receives the mirrored traffic from the source VLAN. (Range: 1-12/28)

Web Interface

To configure VLAN mirroring:

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

Figure 89: Configuring VLAN Mirroring

VLAN > Mirror

Action: Add

Source VLAN: 1

Target Port: 2

Apply Revert

To show the VLANs to be mirrored:

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

Figure 90: Showing the VLANs to Mirror

VLAN > Mirror

Action: Show

VLAN Mirror List Total: 1

	Source (VLAN)	Target (Unit/Port)
<input type="checkbox"/>	1	1/2

Delete Revert

6

Address Table Settings

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

This chapter describes the following topics:

- ◆ [Static MAC Addresses](#) – Configures static entries in the address table.
- ◆ [MAC Address Isolation](#) – Forwards or filters static entries based on specified isolation mode.
- ◆ [Address Aging Time](#) – Sets time out for dynamically learned entries.
- ◆ [Dynamic Address Cache](#) – Shows dynamic entries in the address table.
- ◆ [MAC Address Mirroring](#) – Mirrors traffic matching a specified source address to a target port.
- ◆ [Extended MAC Security](#) – Configures the maximum number of MAC addresses that can be learned on an interface, the movable-static function which allows a static address to be moved to another interface, and the sticky-dynamic function which prevents dynamic address already learned elsewhere from being learned at a specified interface.

Setting Static Addresses

Use the [MAC Address > Static \(Configure MAC Address\)](#) page to configure static MAC addresses. A static address can be assigned to a specific interface on this switch.

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:

- ◆ **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, or to one of the other optional modes.
 - Permanent - Assignment is permanent. (This is the default.)
 - Delete-on-reset - Assignment lasts until the switch is reset.
 - Movable Static - Refer to [“Configuring Extended MAC Security.”](#)
 - Isolation - Refer to [“Configuring MAC Address Isolation.”](#)

Web Interface

To configure a static MAC address:

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

Figure 91: Configuring Static MAC Addresses

MAC Address > Static

Step: 2. Configure MAC Address Action: Add

VLAN: 1

Interface: Port 1 Trunk

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

Static Status: Permanent

Apply Revert

To show the static addresses in MAC address table:

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

Figure 92: Displaying Static MAC Addresses

The screenshot shows the 'MAC Address > Static' configuration page. At the top, there is a breadcrumb 'MAC Address > Static'. Below it, there are two dropdown menus: 'Step: 2. Configure MAC Address' and 'Action: Show'. The main content is a table titled 'Static MAC Address to Interface Mapping Table' with a 'Total: 2' indicator. The table has seven columns: a checkbox, MAC Address, VLAN, Interface, Type, Life Time, and Isolation. There are two rows of data. Below the table are 'Delete' and 'Revert' buttons.

<input type="checkbox"/>	MAC Address	VLAN	Interface	Type	Life Time	Isolation
<input type="checkbox"/>	00-00-00-00-00-02	1	CPU	CPU	Delete on Reset	NA
<input type="checkbox"/>	00-12-CF-94-34-DA	1	Unit 1 / Port 1	Config	Permanent	Promiscuous

Configuring MAC Address Isolation

Use the MAC Address > Static (Configure Global) page to enable MAC address isolation globally on the switch, and the (Configure MAC Address) page to set the isolation mode for a static MAC address.

MAC address isolation provides tighter control over traffic passing from downlink ports over the local network. Access for individual clients can be restricted to uplink ports, or also to members within a shared community.

Command Usage

- ◆ MAC address isolation must be enabled globally on the Configure Global page before isolation profiles defined on the Configure MAC Address – Add page will take effect.
- ◆ MAC address isolation is restricted to static MAC addresses which specify one of the isolation modes.
- ◆ If the MAC address for a packet is found during source address lookup and an isolation profile is assigned to that address, then that profile is used as the source address isolation profile.
- ◆ If the MAC address for a packet is found during destination address lookup and an isolation profile is assigned to that address, then that profile is used as the destination address isolation profile.

- ◆ Packets are filtered or forwarded according to the isolation profiles shown in the following table.

Table 10: MAC Address Isolation Matrix

SA Profile	DA Profile			
	Promiscuous	Isolated	Community 1	Community 2
Promiscuous	Forward	Forward	Forward	Forward
Isolated	Forward	Drop	Drop	Drop
Community 1	Forward	Drop	Forward	Drop
Community 2	Forward	Drop	Drop	Forward

Parameters

These parameters are displayed:

Configure Global

- ◆ **MAC-Isolation Global Status** – Enables MAC isolation globally on the switch. (Default: Disabled)

Configure MAC Address (Static Status - Isolation)

- ◆ **Static Status** – Sets the status to Isolation, or one of the other optional modes.
- ◆ **MAC Isolation Mode** – Sets the MAC address isolation mode to one of the following options:
 - **Community 1** – Forwards packets from a source address defined as community 1 only to a destination address set to promiscuous mode or set to community 1.
 - **Community 2** – Forwards packets from a source address defined as community 2 only to a destination address set to promiscuous mode or set to community 2.
 - **Isolated** – Forwards packets from a source address defined as isolated only to a destination address set to promiscuous mode.
 - **Promiscuous** – Forwards packets from a source address defined as promiscuous to a destination address set to any isolation mode. Note that promiscuous mode should be used for any uplink port that will service downlink ports assigned to one of the other isolation modes.

For a description of the other parameters on this page, refer to [“Setting Static Addresses” on page 165](#).

Web Interface

To enable or disable MAC isolation globally on the switch:

1. Click MAC Address, Static.
2. Select Configure Global from the Step list.
3. Set the MAC-Isolation Global Status.
4. Enter Apply.

Figure 93: Setting Global Status for MAC Address Isolation

MAC Address > Static

Step: 1. Configure Global

MAC-Isolation Global Status Enabled

Apply Revert

To assign static addresses to an isolation mode:

1. Click MAC Address, Static.
2. Select Configure MAC Address from the Step list.
3. Select Add from the Action list.
4. Specify the VLAN, the port or trunk to which the address will be assigned, and the MAC address.
5. Set Static Status to Isolation.
6. Select an option for MAC Isolation Mode.
7. Enter Apply.

Figure 94: Configuring MAC Address Isolation

MAC Address > Static

Step: 2. Configure MAC Address Action: Add

VLAN 1

Interface Port 1 Trunk

MAC Address 00-12-cf-94.33-db

Static Status Isolation

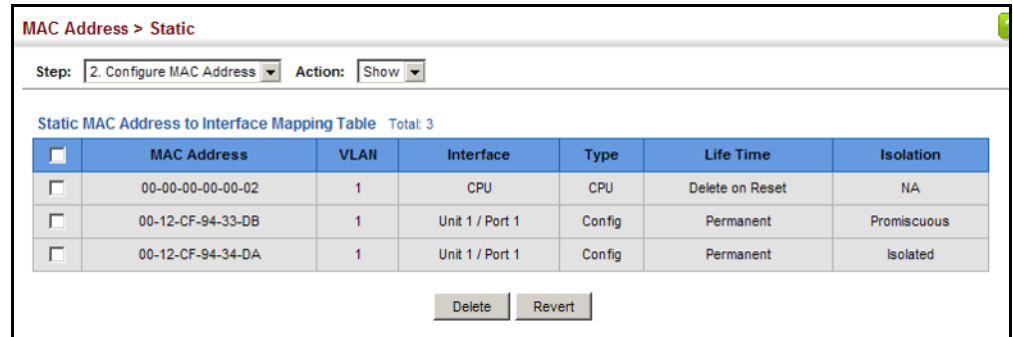
MAC Isolation Mode Promiscuous

Apply Revert

To show the static addresses assigned to an isolation mode:

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

Figure 95: Displaying Interfaces Assigned a MAC Address Isolation Mode



The screenshot shows the 'MAC Address > Static' configuration page. At the top, there is a breadcrumb 'MAC Address > Static'. Below it, the 'Step' is set to '2. Configure MAC Address' and the 'Action' is 'Show'. The main content is a table titled 'Static MAC Address to Interface Mapping Table' with a total of 3 entries. The table has columns for checkboxes, MAC Address, VLAN, Interface, Type, Life Time, and Isolation. Below the table are 'Delete' and 'Revert' buttons.

<input type="checkbox"/>	MAC Address	VLAN	Interface	Type	Life Time	Isolation
<input type="checkbox"/>	00-00-00-00-00-02	1	CPU	CPU	Delete on Reset	NA
<input type="checkbox"/>	00-12-CF-94-33-DB	1	Unit 1 / Port 1	Config	Permanent	Promiscuous
<input type="checkbox"/>	00-12-CF-94-34-DA	1	Unit 1 / Port 1	Config	Permanent	Isolated

Changing the Aging Time

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

Parameters

These parameters are displayed:

- ◆ **Aging Status** – Enables/disables the function.
- ◆ **Aging Time** – The time after which a learned entry is discarded. (Range: 6-7200 seconds; Default: 300 seconds)

Web Interface

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

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

Figure 96: Setting the Address Aging Time

The screenshot shows a web interface for configuring MAC address settings. At the top, it says "MAC Address > Dynamic". Below that, there is a dropdown menu for "Action" set to "Configure Aging". Underneath, "Aging Status" is checked and labeled "Enabled". "Aging Time (6-7200)" is set to "300" with "sec" next to it. At the bottom right, there are "Apply" and "Revert" buttons.

Displaying the Dynamic Address Table

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

Parameters

These parameters are displayed:

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

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 97: Displaying the Dynamic MAC Address Table

The screenshot shows the 'MAC Address > Dynamic' configuration page. At the top, there is an 'Action:' dropdown menu set to 'Show Dynamic MAC'. Below this, the 'Query by:' section includes a 'Sort Key' dropdown set to 'MAC Address'. There are three checkboxes: 'MAC Address' (unchecked), 'VLAN' (unchecked), and 'Interface' (unchecked). The 'VLAN' field is set to '1'. The 'Interface' field has two options: 'Port' (selected) and 'Trunk'. The 'Port' field is set to '1'. A 'Query' button is located below the query fields. Below the query fields, there is a table titled 'Dynamic MAC Address List' with a 'Total: 3' count. The table has seven columns: MAC Address, VLAN, Interface, Type, Life Time, and Isolation. The table contains three rows of data.

MAC Address	VLAN	Interface	Type	Life Time	Isolation
00-E0-0C-10-90-00	1	Unit 1 / Port 1	Learn	Delete on Timeout	NA
00-E0-0C-10-90-0B	1	Unit 1 / Port 1	Learn	Delete on Timeout	NA
00-E0-29-94-34-64	1	Unit 1 / Port 1	Learn	Delete on Timeout	NA

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 98: Clearing Entries in the Dynamic MAC Address Table

The screenshot shows a web interface for clearing dynamic MAC address table entries. At the top, it says 'MAC Address > Dynamic'. Below this, there is a section for 'Action:' with a dropdown menu set to 'Clear Dynamic MAC'. Underneath, there is a 'Clear by:' dropdown menu set to 'All'. A 'Clear' button is located at the bottom right of the form area.

Configuring MAC Address Mirroring

Use the MAC Address > Mirror (Add) page to mirror traffic matching a specified source address from any port on the switch to a target port for real-time analysis. You can then attach a logic analyzer or RMON probe to the target port and study the traffic crossing the source port in a completely unobtrusive manner.

Command Usage

- ◆ When mirroring traffic from a MAC address, ingress traffic with the specified source address entering any port in the switch, other than the target port, will be mirrored to the destination port.
- ◆ All mirror sessions must share the same destination port.
- ◆ Spanning Tree BPDU packets are not mirrored to the target port.
- ◆ When mirroring port traffic, the target port must be included in the same VLAN as the source port when using MSTP (see [“Spanning Tree Algorithm” on page 179](#)).
- ◆ When mirroring VLAN traffic (see [“Configuring VLAN Mirroring” on page 162](#)) or packets based on a source MAC address, the target port cannot be set to the same target ports as that used for port mirroring (see [“Configuring Local Port Mirroring” on page 99](#)).
- ◆ When traffic matches the rules for both port mirroring, and for mirroring of VLAN traffic or packets based on a MAC address, the matching packets will not be sent to target port specified for port mirroring.

Parameters

These parameters are displayed:

- ◆ **Source MAC** – MAC address in the form of xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx.
- ◆ **Target Port** – The port that will mirror the traffic from the source port. (Range: 1-12/28)

Web Interface

To mirror packets based on a MAC address:

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

Figure 99: Mirroring Packets Based on the Source MAC Address

MAC Address > Mirror

Action: Add

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

Target Port: 2

Apply Revert

To show the MAC addresses to be mirrored:

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

Figure 100: Showing the Source MAC Addresses to Mirror

MAC Address > Mirror

Action: Show

MAC Mirror List Total: 1

	Source (MAC)	Target (Unit/Port)
<input type="checkbox"/>	11-22-33-44-55-66	1/2

Delete Revert

Configuring Extended MAC Security

Use the MAC Address > Extended MAC Security pages to configure the maximum number of MAC addresses that can be learned on an interface, the movable-static function which allows a static address to be moved to another interface, and the sticky-dynamic function which prevents dynamic address already learned elsewhere from being learned at a specified interface.

Command Usage

- ◆ At most 32 VLANs can be enabled for the maximum MAC address count, the sticky-dynamic function, and the movable-static function.
- ◆ If a trap is set on an interface and the same security violation occurs, a minimum interval of 60 seconds is used between sending subsequent trap messages to prevent flooding of trap messages.
- ◆ If a trap is set on both a port and the VLAN to which that port belongs, and a security violation occurs, two trap messages will be sent.

Parameters

These parameters are displayed:

- ◆ **VLAN** – VLAN identifier. (Configure VLAN)
- ◆ **Interface** – Port or Trunk identifier. (Configure Interface)
- ◆ **Sticky Dynamic MAC Status** – Prevents source addresses learned at other interfaces from being learned at this interface. (Default: Disabled)
 - Once the sticky-dynamic function is enabled on a interface, the MAC addresses dynamically learned on other interfaces cannot be learned by this interface. If a packet with an address learned on another interface tries to make a station-move to this interface, it will be treated as security breach and discarded.
 - The sticky-dynamic function cannot be set for a port that is a member of a static or dynamic trunk.
- ◆ **Movable Static MAC Status** – Specifies an interface to which a static MAC address can be moved. (Default: Enabled)
 - Use MAC Address > Static (Configure MAC Address - Add) page to add static addresses to the MAC address table. These addresses are not aged out, nor removed from the address table when the assigned interface is down.

When the movable-static address function is enabled on an interface, static MAC addresses can be moved to this interface.

When the movable-static address function is disabled on an interface, and a packet with a static MAC address attempts to enter this interface, the packet is dropped.

- The movable-static function cannot be set for a port that is a member of a static or dynamic trunk. When a trunk is formed, the trunk takes on the movable-static status of the first port to join the trunk. When other ports are subsequently added to a trunk, those ports take on the movable-static status of the trunk. When a port leaves a trunk, it retains the movable-static status of the trunk.
- If port security is enabled (see [“Configuring Port Security” on page 321](#)), it may turn learned MAC addresses into normal static addresses. There is no conflict between this function and the movable-static address function.
- ◆ **Max MAC Count** – Sets the maximum number of MAC addresses which can be learned on an interface. (Range: 0-1024; Default: Disabled)
 - If an interface reaches its maximum learning count, a packet with a new source address will be dropped and address will not be learned.

If the sticky-dynamic function or movable-static function is enabled, and the maximum MAC count has not been set, it will be automatically enabled by the system and set to 1024.

If the maximum MAC count is changed from small count to bigger one, the learned MACs on that interface are all retained. However, if it is changed from a large count to a smaller one, all of the learned MAC addresses on that interface are cleared.
 - The maximum MAC count cannot be set on a trunk, but can set for the member ports of a trunk.
- ◆ **Current MAC Count** – The number of static MAC address entries in the address table for this interface.
- ◆ **Sticky Dynamic MAC Trap Status** – Sends a trap if an ingress packet violates the configured settings for the sticky-dynamic function. (Default: Disabled)
- ◆ **Movable Static MAC Trap Status** – Sends a trap if an ingress packet violates the configured settings for the movable-static function. (Default: Disabled)
- ◆ **Max MAC Count Trap Status** – Sends a trap if an ingress packet violates the configured settings for the maximum MAC address count. (Default: Disabled)

Web Interface

To configure extended MAC security:

1. Click MAC Address, Extended MAC Security.
2. Select Configure VLAN or Configure Interface from the Action list.
3. Select the VLAN, port or trunk to configure.

4. Make the required changes to the sticky-dynamic function, movable-static function, and maximum MAC count. Enable or disable trap messages for these features.
5. Click Apply.

Figure 101: Configuring Extended MAC Security on a VLAN

MAC Address > Extended MAC Security

Step: 1. Configure VLAN

VLAN: 1

Sticky Dynamic MAC Status: Enabled

Movable Static MAC Status: Enabled

Max MAC Count (1-1024): undefined

Current MAC Count: undefined

Sticky Dynamic MAC Trap Status: Enabled

Movable Static MAC Trap Status: Enabled

Max MAC Count Trap Status: Enabled

Apply Revert

Figure 102: Configuring Extended MAC Security on a Port or Trunk

MAC Address > Extended MAC Security

Step: 2. Configure Interface

Interface: Port 1 Trunk

Sticky Dynamic MAC Status: Enabled

Movable Static MAC Status: Enabled

Max MAC Count (1-1024): 0

Current MAC Count: 2

Sticky Dynamic MAC Trap Status: Enabled

Movable Static MAC Trap Status: Enabled

Max MAC Count Trap Status: Enabled

Apply Revert

7

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 103: 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 104: 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 196](#)). 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 105: Common Internal Spanning Tree, Common Spanning Tree, Internal Spanning Tree



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: 30-86400 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.

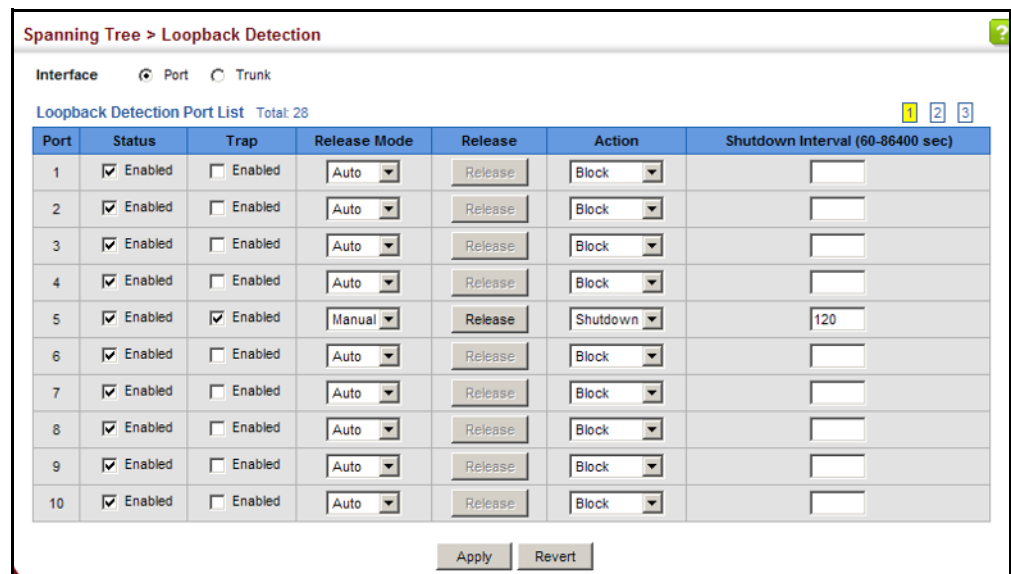
If an interface is shut down due to a detected loopback, and the release mode is set to "Manual," the interface can be re-enabled using the Release button.

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 106: Configuring Port Loopback Detection



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 Protocol²

This option uses RSTP set to STP forced compatibility mode. It uses RSTP for the internal state machine, but sends only 802.1D BPDUs. This creates one spanning tree instance for the entire network. If multiple VLANs are

implemented on a network, the path between specific VLAN members may be inadvertently disabled to prevent network loops, thus isolating group members. When operating multiple VLANs, we recommend selecting the MSTP option.

◆ **Rapid Spanning Tree Protocol²**

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)

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

- ◆ **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**³ – The revision for this MSTI. (Range: 0-65535; Default: 0)
- ◆ **Region Name**³ – The name for this MSTI. (Maximum length: 32 characters; switch’s MAC address)
- ◆ **Max Hop Count** – The maximum number of hops allowed in the MST region before a BPDU is discarded. (Range: 1-40; Default: 20)



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

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

Web Interface

To configure global STA settings:

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

Figure 107: Configuring Global Settings for STA (STP)

Spanning Tree > STA

Step: 1. Configure Global Action: Configure

Spanning Tree Status Enabled

Spanning Tree Type STP

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

BPDU Flooding To VLAN

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 108: 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 109: Configuring Global Settings for STA (MSTP)

Spanning Tree > STA

Step: 1. Configure Global Action: Configure

Spanning Tree Status Enabled

Spanning Tree Type MSTP

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

BPDU Flooding To VLAN

Advanced:

Path Cost Method Long

Transmission Limit (1-10) 3

When the Switch Becomes Root:

Hello Time (1-10) 2 sec

Maximum Age (6-40) 20 sec

Forward Delay (4-30) 15 sec

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

MSTP Configuration

Max Instance Numbers 32

Configuration Digest 0xAC36177F50283CD4B83821D8AB26DE62

Region Revision (0-65535) 0

Region Name 00 e0 0c 00 01 fd

Max Hop Count (1-40) 20

Apply Revert

Displaying Global Settings for STA

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

Parameters

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

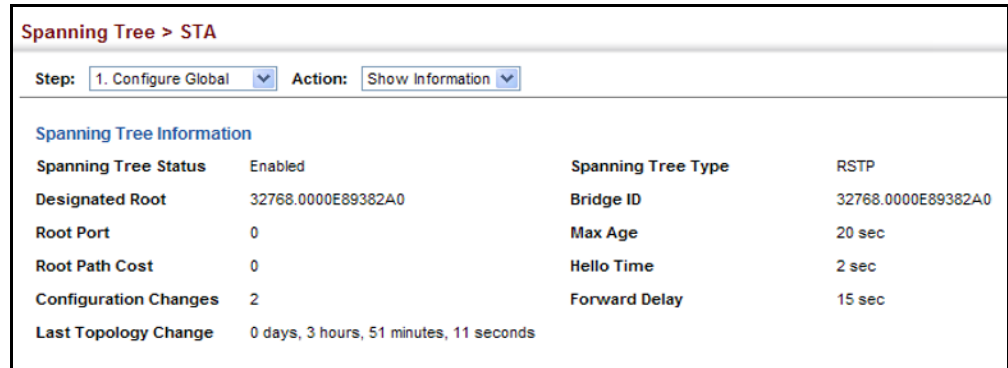
- ◆ **Bridge ID** – A unique identifier for this bridge, consisting of the bridge priority, the MST Instance ID 0 for the Common Spanning Tree when spanning tree type is set to MSTP, and MAC address (where the address is taken from the switch system).
- ◆ **Designated Root** – The priority and MAC address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

Web Interface

To display global STA settings:

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

Figure 110: Displaying Global Settings for STA



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

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 183](#)) or when spanning tree is disabled on a specific port. When flooding is enabled, BPDUs are flooded to all other ports on the switch or to all other ports within the receiving port’s native VLAN as specified by the Spanning Tree BPDU Flooding attribute ([page 183](#)). (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 method⁴, 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

Table 12: Default STA Path Costs

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

- ◆ **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)

4. Refer to “[Configuring Global Settings for STA](#)” on page 183 for information on setting the path cost method.

- ◆ **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 183](#)).

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

- If spanning tree mode is set to STP ([page 183](#)), edge-port mode cannot automatically transition to operational edge-port state using the automatic setting.
 - If loopback detection is enabled ([page 182](#)) 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 194](#)).
- ◆ **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 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)

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

Web Interface

To configure interface settings for STA:

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

Figure 111: Configuring Interface Settings for STA

Spanning Tree > STA

Step: 2. Configure Interface Action: Configure

Interface Port Trunk

Port List Total: 28

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

Displaying Interface Settings for STA

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

Parameters

These parameters are displayed:

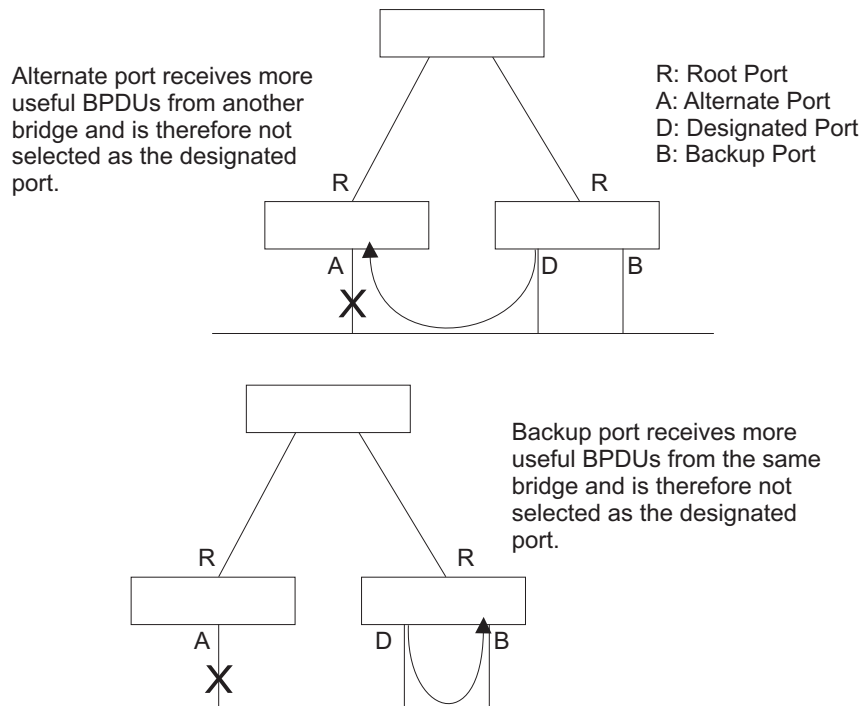
- ◆ **Spanning Tree** – Shows if STA has been enabled on this interface.
- ◆ **BPDU Flooding** – Shows if BPDUs will be flooded to other ports when spanning tree is disabled globally on the switch or disabled on a specific port.
- ◆ **STA Status** – Displays current state of this port within the Spanning 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 190](#).
- ◆ **Oper Edge Port** – This parameter is initialized to the setting for Admin Edge Port in STA Port Configuration on [page 190](#) (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 112: STA Port Roles



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 113: Displaying Interface Settings for STA

The screenshot shows the 'Spanning Tree > STA' configuration page. At the top, there are dropdown menus for 'Step: 2. Configure Interface' and 'Action: Show Information'. Below this, there are radio buttons for 'Interface' with 'Port' selected and 'Trunk' unselected. The main content is a table titled 'Spanning Tree Port List' with a total of 28 ports. The table has 12 columns: Port, Spanning Tree, BPDU Flooding, STA Status, Forward Transitions, Designated Cost, Designated Bridge, Designated Port, Oper Path Cost, Oper Link Type, Oper Edge Port, and Port Role. The first five rows of data are shown below.

Port	Spanning Tree	BPDU Flooding	STA Status	Forward Transitions	Designated Cost	Designated Bridge	Designated Port	Oper Path Cost	Oper Link Type	Oper Edge Port	Port Role
1	Enabled	Enabled	Forwarding	1	100000	32768.00E00C109000	128.2	100000	Point-to-Point	Disabled	Root
2	Enabled	Enabled	Discarding	0	200000	32768.7072CF5BDAD4	128.2	100000	Point-to-Point	Disabled	Disabled
3	Enabled	Enabled	Discarding	0	200000	32768.7072CF5BDAD4	128.3	100000	Point-to-Point	Disabled	Disabled
4	Enabled	Enabled	Discarding	0	200000	32768.7072CF5BDAD4	128.4	100000	Point-to-Point	Disabled	Disabled
5	Enabled	Enabled	Discarding	0	200000	32768.7072CF5BDAD4	128.5	100000	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 32 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 183) 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 183](#)).
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 114: Creating an MST Instance

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

To show the MSTP instances:

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

Figure 115: Displaying MST Instances

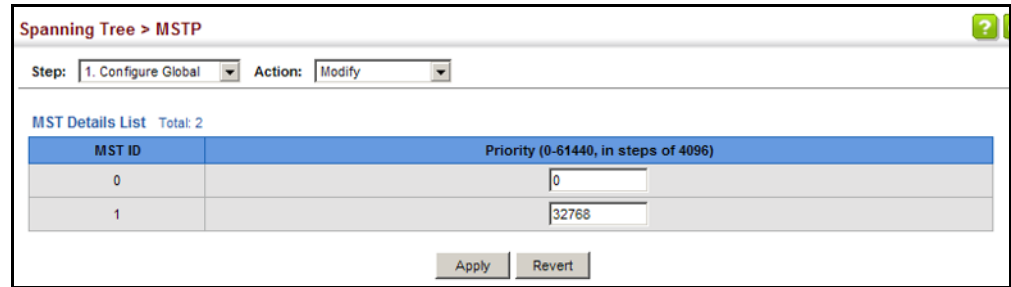
The screenshot shows the 'Spanning Tree > MSTP' configuration page with the 'Action' dropdown set to 'Show'. Below the dropdowns, there is a table titled 'MST List Total: 2'. The table has two columns: a checkbox column and an 'MST ID' column. There are two rows of data: one with MST ID 0 and one with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

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

To modify the priority for an MST instance:

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

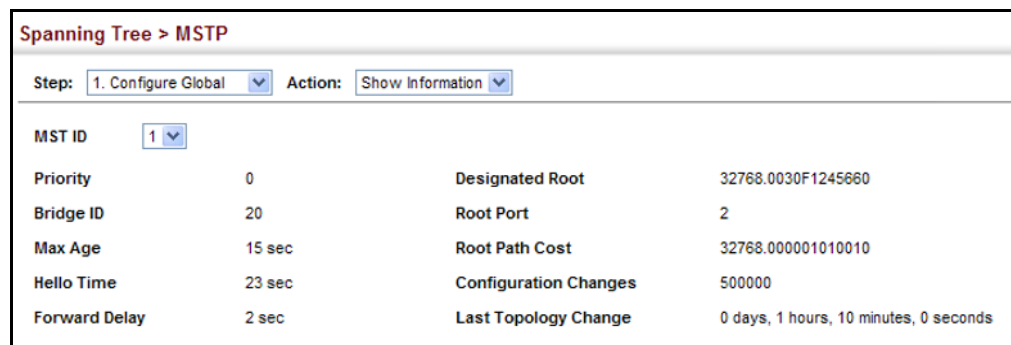
Figure 116: Modifying the Priority for an MST Instance



To display global settings for MSTP:

1. Click Spanning Tree, MSTP.
2. Select Configure Global from the Step list.
3. Select Show Information from the Action list.
4. Select an MST ID. The attributes displayed on this page are described under ["Displaying Global Settings for STA" on page 189.](#)

Figure 117: Displaying Global Settings for an MST Instance



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 118: Adding a VLAN to an MST Instance

Spanning Tree > MSTP

Step: 1. Configure Global Action: Add Member

MST ID 1

VLAN ID (1-4094) 2

Apply Revert

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 119: Displaying Members of an MST Instance

Spanning Tree > MSTP

Step: 1. Configure Global Action: Show Member

MST ID 0

Member List Total: 4094

<input type="checkbox"/>	VLAN
<input type="checkbox"/>	1
<input type="checkbox"/>	2
<input type="checkbox"/>	3
<input type="checkbox"/>	4
<input type="checkbox"/>	5

Configuring Interface Settings for MSTP

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

Parameters

These parameters are displayed:

- ◆ **MST ID** – Instance identifier to configure. (Default: 0)
- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **STA Status** – Displays the current state of this interface within the Spanning Tree. (See “[Displaying Interface Settings for STA](#)” on page 194 for additional information.)

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

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

The recommended range is listed in [Table 11 on page 191](#).
The default path costs are listed in [Table 12 on page 191](#).

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 120: 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 121: 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	2	100000	32768.0.7072CF5BDAD4	128.2	100000	Point-to-Point	Disabled	Root
2	Discarding	0	200000	32768.0.7072CF5BDAD4	128.2	100000	Point-to-Point	Disabled	Disabled
3	Discarding	0	200000	32768.0.7072CF5BDAD4	128.3	100000	Point-to-Point	Disabled	Disabled
4	Discarding	0	200000	32768.0.7072CF5BDAD4	128.4	100000	Point-to-Point	Disabled	Disabled
5	Discarding	0	200000	32768.0.7072CF5BDAD4	128.5	100000	Point-to-Point	Disabled	Disabled

Congestion Control

The switch can set the maximum upload or download data transfer rate for any port. It can 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. 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:

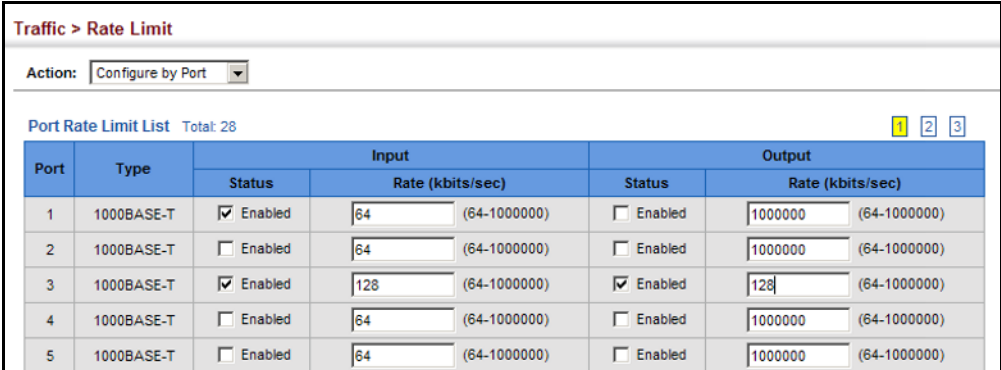
- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Type** – Indicates the port type. (1000BASE-T, 100BASE SFP, or 1000BASE SFP)
- ◆ **Status** – Enables or disables the rate limit. (Default: Disabled)
- ◆ **Rate** – Sets the rate limit level. (Range: 64 - 1,000,000 kbits per second)

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 122: Configuring Rate Limits



The screenshot shows the 'Traffic > Rate Limit' configuration page. At the top, there is a breadcrumb 'Traffic > Rate Limit' and an 'Action:' dropdown menu set to 'Configure by Port'. Below this is a 'Port Rate Limit List' with a 'Total: 28' and three numbered tabs (1, 2, 3). The main table has columns for Port, Type, Input Status, Input Rate (kbits/sec), Output Status, and Output Rate (kbits/sec). The table contains five rows of data for 1000BASE-T ports.

Port	Type	Input		Output	
		Status	Rate (kbits/sec)	Status	Rate (kbits/sec)
1	1000BASE-T	<input checked="" type="checkbox"/> Enabled	64 (64-1000000)	<input type="checkbox"/> Enabled	1000000 (64-1000000)
2	1000BASE-T	<input type="checkbox"/> Enabled	64 (64-1000000)	<input type="checkbox"/> Enabled	1000000 (64-1000000)
3	1000BASE-T	<input checked="" type="checkbox"/> Enabled	128 (64-1000000)	<input checked="" type="checkbox"/> Enabled	128 (64-1000000)
4	1000BASE-T	<input type="checkbox"/> Enabled	64 (64-1000000)	<input type="checkbox"/> Enabled	1000000 (64-1000000)
5	1000BASE-T	<input type="checkbox"/> Enabled	64 (64-1000000)	<input type="checkbox"/> Enabled	1000000 (64-1000000)

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

- ◆ Storm Control is disabled by default.
- ◆ Broadcast control does not effect IP multicast traffic.
- ◆ 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.
- ◆ Storm control is a hardware level function. Traffic storms can also be controlled at the software level using automatic storm control which triggers various control responses (see [“Automatic Traffic Control” on page 206](#)). 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 this 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. For example, suppose broadcast storm control is set to 5000 Kbps, and the rate limit is set to 100000 Kbps on a Gigabit Ethernet port. Since 200000 Kbps is 1/5 of line speed, the received rate will actually be 1000 Kbps, or 1/5 of the 5000 Kbps limit set by the storm control command. It is therefore not advisable to use both of these commands on the same interface.
- ◆ The description of effective rate limiting (see Command Usage under [“Rate Limiting” on page 203](#)) also applies to storm control.

Parameters

These parameters are displayed:

- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **Type** – Indicates interface type. (1000BASE-T or 1000BASE 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., kilobits per second. (Range: 64-1000000 Kbps)

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 123: Configuring Storm Control

Traffic > Storm Control								
Interface <input checked="" type="radio"/> Port <input type="radio"/> Trunk								
Port Storm Control List Total: 28 1 2 3								
Port	Type	Unknown Unicast		Multicast		Broadcast		
		Status	Rate (kbits/sec)	Status	Rate (kbits/sec)	Status	Rate (kbits/sec)	
1	1000BASE-T	<input checked="" type="checkbox"/> Enabled	128 (64-10000000)	<input type="checkbox"/> Enabled	64 (64-10000000)	<input checked="" type="checkbox"/> Enabled	128 (64-10000000)	
2	1000BASE-T	<input type="checkbox"/> Enabled	64 (64-10000000)	<input type="checkbox"/> Enabled	64 (64-10000000)	<input checked="" type="checkbox"/> Enabled	64 (64-10000000)	
3	1000BASE-T	<input checked="" type="checkbox"/> Enabled	600000 (64-10000000)	<input checked="" type="checkbox"/> Enabled	600000 (64-10000000)	<input checked="" type="checkbox"/> Enabled	600000 (64-10000000)	
4	1000BASE-T	<input type="checkbox"/> Enabled	64 (64-10000000)	<input type="checkbox"/> Enabled	64 (64-10000000)	<input checked="" type="checkbox"/> Enabled	64 (64-10000000)	
5	1000BASE-T	<input type="checkbox"/> Enabled	64 (64-10000000)	<input type="checkbox"/> Enabled	64 (64-10000000)	<input checked="" type="checkbox"/> Enabled	64 (64-10000000)	

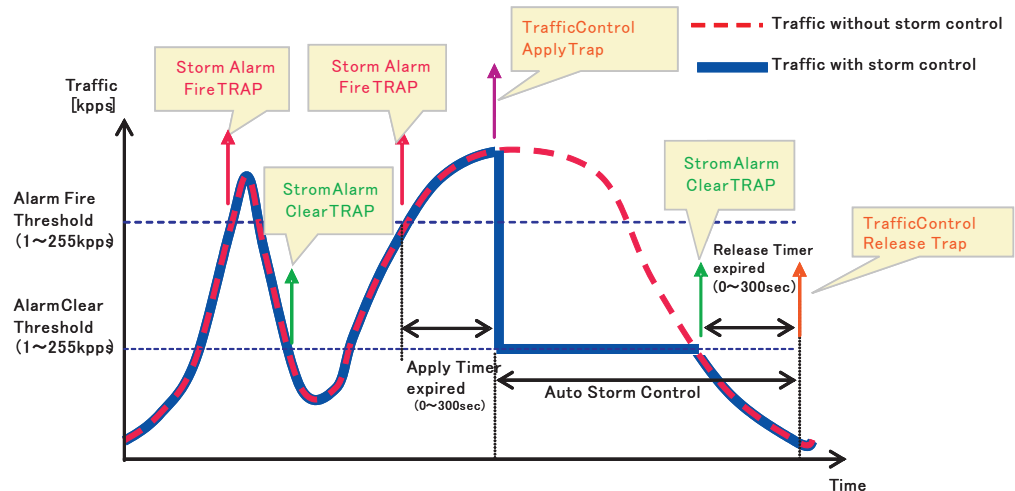
Automatic Traffic Control

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

Command Usage

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

Figure 124: Storm Control by Limiting the Traffic Rate

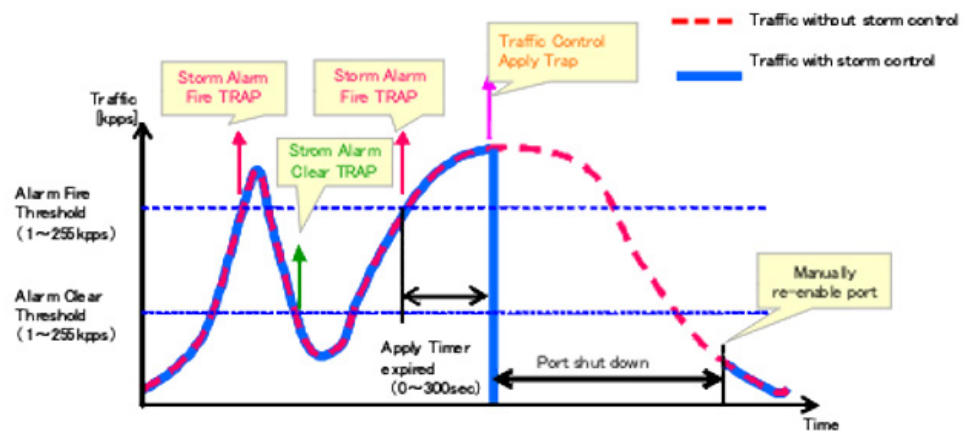


The key elements of this diagram are described below:

- ◆ 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 209](#)).
- ◆ 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 125: 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 204](#)). 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 > Congestion Control > 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 209](#)).

Parameters

These parameters are displayed:

- ◆ **Broadcast Apply Timer** – The interval after the upper threshold has been exceeded at which to apply the control response to broadcast storms. (Range: 1-300 seconds; Default: 300 seconds)
- ◆ **Broadcast Release Timer** – The time at which to release the control response after ingress traffic has fallen beneath the lower threshold for broadcast storms. (Range: 1-900 seconds; Default: 900 seconds)
- ◆ **Multicast Apply Timer** – The interval after the upper threshold has been exceeded at which to apply the control response to multicast storms. (Range: 1-300 seconds; Default: 300 seconds)
- ◆ **Multicast Release Timer** – The time at which to release the control response after ingress traffic has fallen beneath the lower threshold for multicast storms. (Range: 1-900 seconds; Default: 900 seconds)

Web Interface

To configure the response timers for automatic storm control:

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

Figure 126: Configuring ATC Timers

The screenshot shows the configuration interface for Automatic Traffic Control (ATC) timers. The breadcrumb path is 'Traffic > Auto Traffic Control'. The current step is '1. Configure Global'. There are four input fields for timer values, each followed by a 'sec' label. The values are: Broadcast Apply Timer (1-300) is 300, Broadcast Release Timer (1-900) is 900, Multicast Apply Timer (1-300) is 300, and Multicast Release Timer (1-900) is 900. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

Configuring ATC Thresholds and Responses

Use the Traffic > Congestion Control > 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.
Automatic storm control can be enabled for either broadcast or multicast traffic. It cannot be enabled for both of these traffic types at the same time.
- ◆ **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 124 on page 206](#). 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.

If this function is enabled for any port, clicking Apply will manually release the control response, and clear the check box.

Web Interface

To configure the response timers for automatic storm control:

1. Click Traffic, Congestion Control, Automatic Storm 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 127: Configuring ATC Interface Attributes

Traffic > Auto Traffic Control ?

Step: 2. Configure Interface

Storm Control Broadcast Multicast

Auto Traffic Control Broadcast List Total: 28 1 2 3

Port	State	Action	Auto Release Control	Alarm Fire Threshold (1-255 kpps)	Alarm Clear Threshold (1-255 kpps)	Trap Storm Fire	Trap Storm Clear	Trap Traffic Apply	Trap Traffic Release	Manual Control Release
1	<input type="checkbox"/> Enabled	Rate Control	<input type="checkbox"/> Enabled	128	128	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled
2	<input type="checkbox"/> Enabled	Rate Control	<input type="checkbox"/> Enabled	128	128	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled
3	<input type="checkbox"/> Enabled	Rate Control	<input type="checkbox"/> Enabled	128	128	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled
4	<input type="checkbox"/> Enabled	Rate Control	<input type="checkbox"/> Enabled	128	128	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled
5	<input type="checkbox"/> Enabled	Rate Control	<input type="checkbox"/> Enabled	128	128	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled

Class of Service

Class of Service (CoS) allows you to specify which data packets have greater precedence when traffic is buffered in the switch due to congestion. This switch supports CoS with four 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:

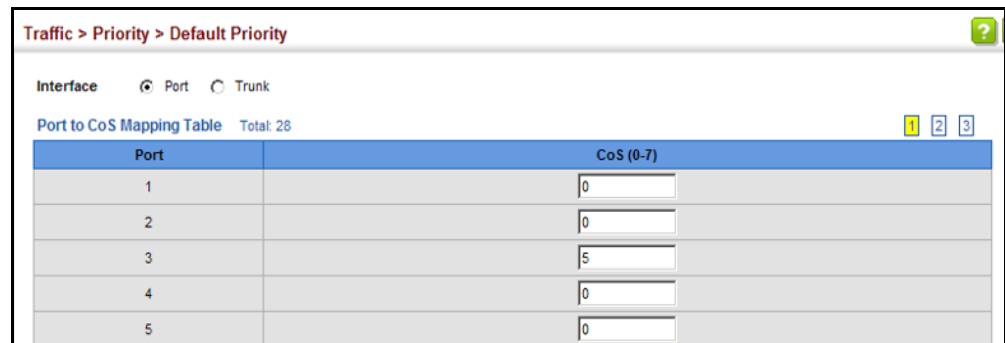
- ◆ **Interface** – Displays a list of ports or trunks.
- ◆ **CoS** – The priority that is assigned to untagged frames received on the specified interface. (Range: 0-7; Default: 0)

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 128: Setting the Default Port Priority



The screenshot shows a web interface for configuring network settings. The breadcrumb trail is 'Traffic > Priority > Default Priority'. There are radio buttons for 'Port' (selected) and 'Trunk'. Below this is a 'Port to CoS Mapping Table' with a 'Total: 28' and three numbered tabs (1, 2, 3). The table has two columns: 'Port' and 'CoS (0-7)'. The 'CoS' column contains input fields with the following values: 0, 0, 5, 0, 0.

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

Service time is shared at the egress ports by defining scheduling weights for WRR, or the queuing 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:

- ◆ **Queue Mode**
 - **Strict** – Services the egress queues in sequential order, transmitting all traffic in the higher priority queues before servicing lower priority queues. This ensures that the highest priority packets are always serviced first, ahead of all other traffic.
 - **WRR** – Weighted Round-Robin shares bandwidth at the egress ports by using scheduling weights, and servicing each queue in a round-robin fashion. (This is the default setting.)
 - **Strict and WRR** – Uses strict or weighted service as specified for each queue.
- ◆ **Queue ID** – The ID of the priority queue. (Range: 0-7)
- ◆ **Strict Mode** – If “Strict and WRR” mode is selected, then a combination of strict and weighted service is used as specified for each queue. Use this parameter to specify the queues assigned to use strict priority when using the strict-weighted queuing mode. (Default: Disabled)
- ◆ **Weight** – Sets a weight for each queue which is used by the WRR scheduler. (Range: 1-255; Default: Weights 1, 2, 4, 6, 8, 10, 12 and 14 are assigned to queues 0 - 7 respectively)

Web Interface

To configure the queue mode:

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

Figure 129: Setting the Queue Mode (Strict)

The screenshot shows the web interface for configuring queue settings. The breadcrumb navigation is "Traffic > Priority > Queue". The "Queue Mode" dropdown menu is set to "Strict". There are "Apply" and "Revert" buttons at the bottom right.

Figure 130: Setting the Queue Mode (WRR)

The screenshot shows the web interface for configuring queue settings. The breadcrumb navigation is "Traffic > Priority > Queue". The "Queue Mode" dropdown menu is set to "WRR". Below the dropdown is a "Queue Setting Table" with a "Total: 8" label. The table has two columns: "Queue ID" and "Weight (1-255)". The table contains 8 rows with Queue IDs from 0 to 7 and corresponding weights from 1 to 14. There are "Apply" and "Revert" buttons at the bottom right.

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

Figure 131: Setting the Queue Mode (Strict and WRR)

Traffic > Priority > Queue

Queue Mode:

Queue Setting Table Total: 8

Queue ID	Strict Mode	Weight (1-255)
0	<input type="text" value="Disabled"/>	<input type="text" value="1"/>
1	<input type="text" value="Disabled"/>	<input type="text" value="2"/>
2	<input type="text" value="Disabled"/>	<input type="text" value="4"/>
3	<input type="text" value="Disabled"/>	<input type="text" value="6"/>
4	<input type="text" value="Disabled"/>	<input type="text" value="8"/>
5	<input type="text" value="Disabled"/>	<input type="text" value="10"/>
6	<input type="text" value="Disabled"/>	<input type="text" value="12"/>
7	<input type="text" value="Disabled"/>	<input type="text" value="14"/>

Mapping CoS Values to Egress Queues

Use the Traffic > Priority > PHB to Queue page to specify the hardware output queues to use based on the internal per-hop behavior value. (For more information on exact manner in which the ingress priority tags are mapped to egress queues for internal processing, see [“Mapping CoS Priorities to Internal DSCP Values”](#) on page 224).

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 224).

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

Parameters

These parameters are displayed:

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

Web Interface

To map internal PHB to hardware queues:

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

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

Figure 132: 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.

Figure 133: Showing CoS Values to Egress Queue Mapping

Traffic > Priority > PHB to Queue

Action: Show

Port: 1

PHB to Queue Mapping List Total: 8

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

Default Revert

Layer 3/4 Priority Settings

Mapping Layer 3/4 Priorities to CoS Values

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

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



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

Setting Priority Processing to DSCP or CoS

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

Command Usage

- ◆ If the QoS mapping mode is set to DSCP, and the ingress packet type is IPv4, then priority processing will be based on the DSCP value in the ingress packet.
- ◆ If the QoS mapping mode is set to DSCP, and a non-IP packet is received, the packet's CoS and CFI (Canonical Format Indicator) values are used for priority processing if the packet is tagged. For an untagged packet, the default port priority (see [page 213](#)) 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 213](#)) is used for priority processing.

Parameters

These parameters are displayed:

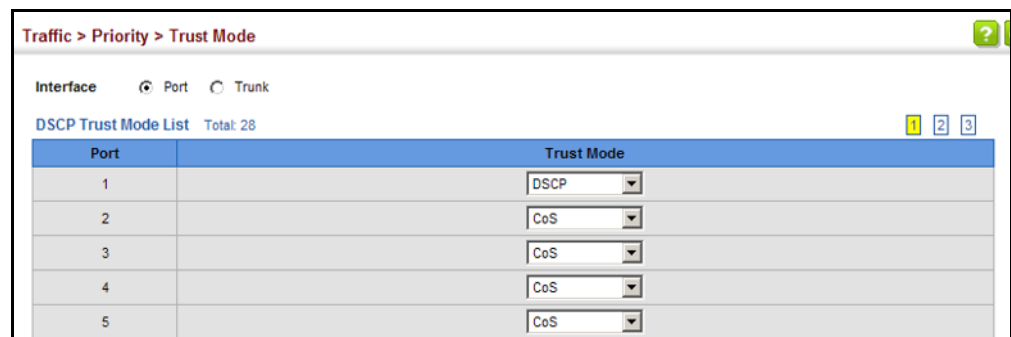
- ◆ **Interface** – Specifies a port or trunk.
- ◆ **Trust Mode**
 - **CoS** – Maps layer 3/4 priorities using Class of Service values. (This is the default setting.)
 - **DSCP** – Maps layer 3/4 priorities using Differentiated Services Code Point values.

Web Interface

To configure the trust mode:

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

Figure 134: Setting the Trust Mode



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 220](#)), 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 in controlling traffic congestion. (Range: 0 - Green, 3 - Yellow, 1 - Red)

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

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

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

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

Web Interface

To map DSCP values to internal PHB/drop precedence:

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

Figure 135: Configuring DSCP to DSCP Internal Mapping

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

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

Figure 136: Showing DSCP to DSCP Internal Mapping

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

Mapping CoS Priorities to Internal DSCP Values

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

Command Usage

- ◆ The default mapping of CoS to PHB values is shown in [Table 17 on page 224](#).
- ◆ Enter up to eight CoS/CFI paired values, per-hop behavior and drop precedence.
- ◆ If a packet arrives with a 802.1Q header but it is not an IP packet, then the CoS/CFI-to-PHB/Drop Precedence mapping table is used to generate priority and drop precedence values for internal processing. Note that priority tags in the original packet are not modified by this command.
- ◆ The internal DSCP consists of three bits for per-hop behavior (PHB) which determines the queue to which a packet is sent; and two bits for drop precedence (namely color) which is used to control traffic congestion.

Parameters

These parameters are displayed:

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

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

CoS	CFI	0	1
0		(0,0)	(0,0)
1		(1,0)	(1,0)
2		(2,0)	(2,0)
3		(3,0)	(3,0)
4		(4,0)	(4,0)
5		(5,0)	(5,0)
6		(6,0)	(6,0)
7		(7,0)	(7,0)

Web Interface

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

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

Figure 137: Configuring CoS to DSCP Internal Mapping

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 138: Showing CoS to DSCP Internal Mapping

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

Quality of Service

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

- ◆ **Class Map** – Creates a map which identifies a specific class of traffic.
- ◆ **Policy Map** – Sets the boundary parameters used for monitoring inbound traffic, and the action to take for conforming and non-conforming traffic.
- ◆ **Binding to a Port** – Applies a policy map to an ingress port.

Overview

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

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

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



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

Note: You should create a class map before creating a policy map. Otherwise, you will not be able to select a class map from the policy rule settings screen (see [page 232](#)).

Command Usage

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

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

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 232](#)) to create a service policy ([page 241](#)) for a specific interface that defines packet classification, service tagging, and bandwidth policing. Note that one or more class maps can be assigned to a policy map.
- ◆ Up to 32 class maps can be configured.

Parameters

These parameters are displayed:

Add

- ◆ **Class Name** – Name of the class map. (Range: 1-32 characters)
- ◆ **Type** – The criteria specified by the match command.
 - **Match All** – Match all conditions within a class map.
 - **Match Any** – Match any condition within a class map.
- ◆ **Description** – A brief description of a class map. (Range: 1-64 characters)

Add Rule

- ◆ **Class Name** – Name of the class map.
- ◆ **Type** – The criteria specified by the match command. (This field is set on the Add page.)
- ◆ **ACL** – Name of an access control list. Any type of ACL can be specified, including standard or extended IPv4/IPv6 ACLs and MAC ACLs.
- ◆ **IP DSCP** – A DSCP value. (Range: 0-63)
- ◆ **IP Precedence** – An IP Precedence value. (Range: 0-7)
- ◆ **IPv6 DSCP** – A DSCP value contained in an IPv6 packet. (Range: 0-63)
- ◆ **VLAN ID** – A VLAN identifier. (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. Set the Type to match any rule or all rules.
6. Enter a description.
7. Click Add.

Figure 139: Configuring a Class Map

Traffic > DiffServ

Step: 1. Configure Class Action: Add

Class Name: rd-class

Type: Match Any

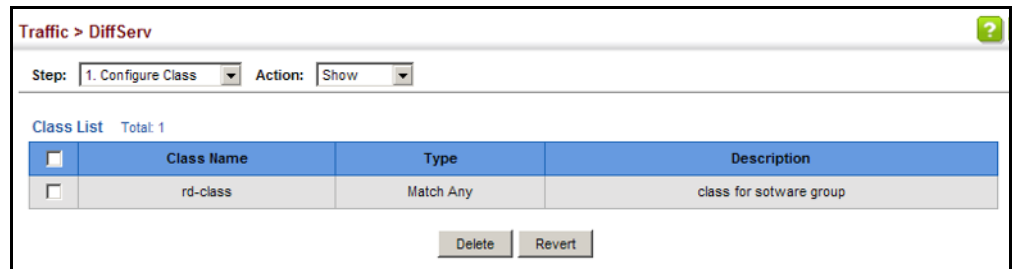
Description: class for software group

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 140: Showing Class Maps



To edit the rules for a class map:

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

Figure 141: Adding Rules to a Class Map

The screenshot shows the configuration interface for DiffServ. The breadcrumb is 'Traffic > DiffServ'. The 'Step' dropdown is set to '1. Configure Class' and the 'Action' dropdown is set to 'Add Rule'. The 'Class Name' is 'rd-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)'. The 'IP DSCP (0-63)' option is selected, and its corresponding input field contains the value '3'. 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 142: Showing the Rules for a Class Map

The screenshot shows the configuration interface for DiffServ. The breadcrumb is 'Traffic > DiffServ'. The 'Step' dropdown is set to '1. Configure Class' and the 'Action' dropdown is set to 'Show Rule'. The 'Class Name' is 'rd-class' and the 'Type' is 'Match Any'. Below this, there is a 'Rule List' section with a 'Total: 2' indicator. A table displays the rules:

<input type="checkbox"/>	Rule
<input type="checkbox"/>	IP DSCP 3
<input type="checkbox"/>	IP Precedence 3

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

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 228), modify service tagging, and enforce bandwidth policing. A policy map can then be bound by a service policy to one or more interfaces (page 241).

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 marker (re)colors an IP packet according to the results of the meter. The color is coded in the DS field [RFC 2474] of the packet.
- ◆ The behavior of the meter is specified in terms of its mode and two token buckets, P and C, which are based on the rates PIR and CIR, respectively. The maximum size of the token bucket P is BP and the maximum size of the token bucket C is BC.

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

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

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

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

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

which are green, yellow, or red. Refer to RFC 2698 for more information on other aspects of trTCM.

Command Usage

- ◆ A policy map can contain 128 class statements that can be applied to the same interface (page 241). 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 241) to take effect.

Parameters

These parameters are displayed:

Add

- ◆ **Policy Name** – Name of policy map. (Range: 1-32 characters)
- ◆ **Description** – A brief description of a policy map. (Range: 1-64 characters)

Add Rule

- ◆ **Policy Name** – Name of policy map.
- ◆ **Class Name** – Name of a class map that defines a traffic classification upon which a policy can act.
- ◆ **Action** – This attribute is used to set an internal QoS value in hardware for matching packets. The PHB label is composed of five bits, three bits for per-hop behavior, and two bits for the color scheme used to control queue congestion with the srTCM and trTCM metering functions.
 - **Set CoS** – Configures the service provided to ingress traffic by setting an internal CoS value for a matching packet (as specified in rule settings for a class map). (Range: 0-7)
See Table 17, “Default Mapping of CoS/CFI to Internal PHB/Drop Precedence,” on page 224).
 - **PHB** – Per-hop behavior, or the priority used for this router hop. (Range: 0-7)
 - **Drop Precedence** – Drop precedence used in controlling traffic congestion. (Range: 0 - Green, 3 - Yellow, 1 - Red)
 - **Set PHB** – Configures the service provided to ingress traffic by setting the internal per-hop behavior for a matching packet (as specified in rule settings for a class map). (Range: 0-7)
See Table 16, “Default Mapping of DSCP Values to Internal PHB/Drop Values,” on page 222).

- **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)
 - **PHB** – Per-hop behavior, or the priority used for this router hop. (Range: 0-7)
 - **Drop Precedence** – Drop precedence used in controlling traffic congestion. (Range: 0 - Green, 3 - Yellow, 1 - Red)
- ◆ **Meter** – Check this to define the maximum throughput, burst rate, and the action that results from a policy violation.
- ◆ **Meter Mode** – Selects one of the following policing methods.
 - **Flow** (Police Flow) – Defines the committed information rate (CIR, or maximum throughput), committed burst size (BC, or burst rate), and the action to take for conforming and non-conforming traffic. Policing is based on a token bucket, where bucket depth (that is, the maximum burst before the bucket overflows) is specified by the “burst” field, and the average rate tokens are removed from the bucket is by specified by the “rate” option.
 - **Committed Information Rate** (CIR) – Rate in kilobits per second. (Range: 0-1000000 kbps at a granularity of 64 kbps or maximum port speed, whichever is lower)
The rate cannot exceed the configured interface speed.
 - **Committed Burst Size** (BC) – Burst in bytes. (Range: 1-2147000 bytes at a granularity of 4k bytes)
 - **Conform** – Specifies that traffic conforming to the maximum rate (CIR) will be transmitted without any change to the DSCP service level.
 - **Transmit** – Transmits in-conformance traffic without any change to the DSCP service level.
 - **Violate** – Specifies whether the traffic that exceeds the maximum rate (CIR) will be dropped or the DSCP service level will be reduced.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63)
 - **Drop** – Drops out of conformance traffic.
 - **srTCM** (Police Meter) – Defines the committed information rate (CIR, or maximum throughput), committed burst size (BC, or burst rate) and excess burst size (BE), and the action to take for traffic conforming to the maximum throughput, exceeding the maximum throughput but within the excess burst size, or exceeding the excess burst size. In addition to the actions defined by this command to transmit, remark the DSCP service value, or drop a packet, the switch will also mark the two color bits used to set the drop precedence of a packet.

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

- **Committed Information Rate (CIR)** – Rate in kilobits per second. (Range: 0-1000000 kbps at a granularity of 64 kbps or maximum port speed, whichever is lower)
The rate cannot exceed the configured interface speed.
- **Committed Burst Size (BC)** – Burst in bytes. (Range: 1-2147000 bytes at a granularity of 4k bytes)
- **Excess Burst Size (BE)** – Burst in excess of committed burst size. (Range: 1-2147000 bytes at a granularity of 4k bytes)
- **Conform** – Specifies that traffic conforming to the maximum rate (CIR) will be transmitted without any change to the DSCP service level.
 - **Transmit** – Transmits in-conformance traffic without any change to the DSCP service level.
- **Exceed** – Specifies whether traffic that exceeds the maximum rate (CIR) but is within the excess burst size (BE) will be dropped or the DSCP service level will be reduced.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63)
 - **Drop** – Drops out of conformance traffic.
- **Violate** – Specifies whether the traffic that exceeds the excess burst size (BE) will be dropped or the DSCP service level will be reduced.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63)
 - **Drop** – Drops out of conformance traffic.
- **trTCM (Police Meter)** – Defines the committed information rate (CIR, or maximum throughput), peak information rate (PIR), and their associated burst sizes – committed burst size (BC, or burst rate) and peak burst size (BP), and the action to take for traffic conforming to the maximum throughput, exceeding the maximum throughput but within the peak information rate, or exceeding the peak information rate. In addition to the actions defined by this command to transmit, remark the DSCP service value, or drop a packet, the switch will also mark the two color bits used to set the drop precedence of a packet.

The color modes include “Color-Blind” which assumes that the packet stream is uncolored, and “Color-Aware” which assumes that the incoming

packets are pre-colored. The functional differences between these modes is described at the beginning of this section under “trTCM Police Meter.”

- **Committed Information Rate (CIR)** – Rate in kilobits per second.
(Range: 0-1000000 kbps at a granularity of 64 kbps or maximum port speed, whichever is lower)

The rate cannot exceed the configured interface speed.
- **Committed Burst Size (BC)** – Burst in bytes. (Range: 1-2147000 bytes at a granularity of 4k bytes)
- **Peak Information Rate (PIR)** – Rate in kilobits per second.
(Range: 0-1000000 kbps at a granularity of 64 kbps or maximum port speed, whichever is lower)

The rate cannot exceed the configured interface speed.
- **Peak Burst Size (BP)** – Burst size in bytes. (Range: 1-2147000 bytes at a granularity of 4k bytes)
- **Conform** – Specifies that traffic conforming to the maximum rate (CIR) will be transmitted without any change to the DSCP service level.
 - **Transmit** – Transmits in-conformance traffic without any change to the DSCP service level.
- **Exceed** – Specifies whether traffic that exceeds the maximum rate (CIR) but is within the peak information rate (PIR) will be dropped or the DSCP service level will be reduced.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63).
 - **Drop** – Drops out of conformance traffic.
- **Violate** – Specifies whether the traffic that exceeds the peak information rate (PIR) will be dropped or the DSCP service level will be reduced.
 - **Set IP DSCP** – Decreases DSCP priority for out of conformance traffic. (Range: 0-63).
 - **Drop** – Drops out of conformance traffic.

Web Interface

To configure a policy map:

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

3. Select Add from the Action list.
4. Enter a policy name.
5. Enter a description.
6. Click Add.

Figure 143: 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 144: Showing Policy Maps

Traffic > DiffServ

Step: 2. Configure Policy Action: Show

Policy List Total: 1

<input type="checkbox"/>	Policy Name	Description
<input type="checkbox"/>	rd-policy	for the software group

Delete Revert

To edit the rules for a policy map:

1. Click Traffic, DiffServ.
2. Select Configure Policy from the Step list.
3. Select Add Rule from the Action list.
4. Select the name of a policy map.
5. Set the CoS or per-hop behavior for matching packets to specify the quality of service to be assigned to the matching traffic class. Use one of the metering

options to define parameters such as the maximum throughput and burst rate. Then specify the action to take for conforming traffic, the action to take for traffic in excess of the maximum rate but within the peak information rate, or the action to take for a policy violation.

6. Click Apply.

Figure 145: Adding Rules to a Policy Map

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 146: Showing the Rules for a Policy Map

			Meter							
<input type="checkbox"/>	Class Name	Action	Meter Mode	Committed Information Rate (kbps)	Committed Burst Size (kbits)	Excess Information Rate (kbps)	Excess Burst Size (kbits)	Conform	Exceed	Violate
<input type="checkbox"/>	rd-class	Set CoS 3, PHB 3, DP 0	Flow	1000000	4000					Drop

Attaching a Policy Map to a Port

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

Command Usage

First define a class map, define a policy map, and then bind the service policy to the required interface.

Parameters

These parameters are displayed:

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

Web Interface

To bind a policy map to a port:

1. Click Traffic, DiffServ.
2. Select Configure Interface from the Step list.
3. Check the box under the Ingress or Egress field to enable a policy map for a port.
4. Select a policy map from the scroll-down box.
5. Click Apply.

Figure 147: Attaching a Policy Map to a Port

Traffic > DiffServ

Step: 3. Configure Interface

Port Service Policy List Total: 28 1 2 3

Port	Ingress	Egress
1	<input type="checkbox"/> accounting	<input type="checkbox"/> accounting
2	<input checked="" type="checkbox"/> accounting	<input checked="" type="checkbox"/> rd-policy
3	<input type="checkbox"/> accounting	<input type="checkbox"/> accounting
4	<input type="checkbox"/> accounting	<input type="checkbox"/> accounting
5	<input type="checkbox"/> accounting	<input type="checkbox"/> accounting
6	<input type="checkbox"/> accounting	<input type="checkbox"/> accounting
7	<input type="checkbox"/> accounting	<input type="checkbox"/> accounting
8	<input type="checkbox"/> accounting	<input type="checkbox"/> accounting
9	<input type="checkbox"/> accounting	<input type="checkbox"/> accounting
10	<input type="checkbox"/> accounting	<input type="checkbox"/> accounting

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 access 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 140](#)).

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 148: 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. Enter the MAC address in format 01-23-45-67-89-AB.
- ◆ **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. (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 149: 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 150: Showing an OUI Telephony List

Traffic > VoIP

Step: 2. Configure OUI Action: Show

Telephony OUI List Total: 3

<input type="checkbox"/>	Telephony OUI	Mask	Description
<input type="checkbox"/>	00-E0-BB-00-00-00	FF-FF-FF-00-00-00	old phones
<input type="checkbox"/>	00-11-22-33-44-55	FF-FF-FF-00-00-00	new phones
<input type="checkbox"/>	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 access 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 140).

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 359](#) 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.

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 151: 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	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> OUI <input type="checkbox"/> LLDP	6	NA
2	None	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> OUI <input type="checkbox"/> LLDP	6	NA
3	Manual	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> OUI <input type="checkbox"/> LLDP	5	NA
4	None	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> OUI <input type="checkbox"/> LLDP	6	NA
5	None	<input type="checkbox"/> Enabled	<input checked="" type="checkbox"/> OUI <input type="checkbox"/> LLDP	6	NA

Security Measures

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

- ◆ **AAA** – Use local or remote authentication to configure access rights, specify authentication servers, configure remote authentication and accounting.
- ◆ **User Accounts** – Manually configure access rights on the switch for specified users.
- ◆ **Web Authentication** – Allows stations to authenticate and access the network in situations where 802.1X or Network Access authentication methods are infeasible or impractical.
- ◆ **Network Access** - Configure MAC authentication, intrusion response, dynamic VLAN assignment, and dynamic QoS assignment.
- ◆ **HTTPS** – Provide a secure web connection.
- ◆ **SSH** – Provide a secure shell (for secure Telnet access).
- ◆ **ACL** – Access Control Lists provide packet filtering for IP frames (based on address, protocol, Layer 4 protocol port number or TCP control code).
- ◆ **ARP Inspection** – Security feature that validates the MAC Address bindings for Address Resolution Protocol packets. Provides protection against ARP traffic with invalid MAC to IP Address bindings, which forms the basis for certain “man-in-the-middle” attacks.
- ◆ **IP Filter** – Filters management access to the web, SNMP or Telnet interface.
- ◆ **Port Security** – Configure secure addresses for individual ports.
- ◆ **Port Authentication** – Use IEEE 802.1X port authentication to control access to specific ports.
- ◆ **DoS Protection** – Protects against Denial-of-Service attacks.
- ◆ **IP Source Guard** – Filters untrusted DHCP messages on insecure ports by building and maintaining a DHCP snooping binding table.
- ◆ **DHCP Snooping** – Filter IP traffic on insecure ports for which the source address cannot be identified via DHCP snooping.



Note: The priority of execution for the filtering commands is Port Security, Port Authentication, Network Access, Web Authentication, Access Control Lists, IP Source Guard, and then DHCP Snooping.

AAA 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 251](#).
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 152: Configuring the Authentication Sequence

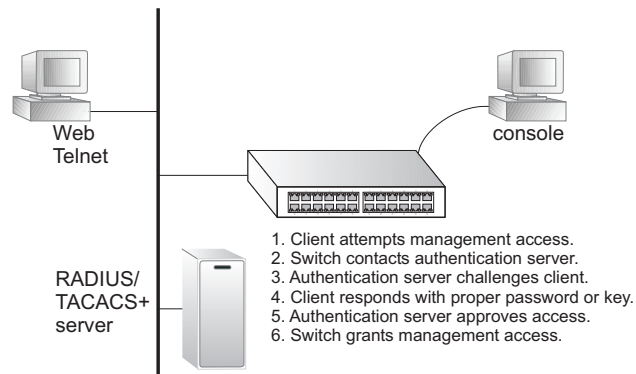


Configuring Remote Logon Authentication Servers

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 153: 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. Do not use blank spaces in the string. (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. Do not use blank spaces in the string. (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 251](#)).

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 154: Configuring Remote Authentication Server (RADIUS)

The screenshot shows the configuration page for a RADIUS server. The breadcrumb navigation is "Security > AAA > Server". The "Server Type" is set to "RADIUS" (selected with a radio button) and "TACACS+" is unselected. Under "Server Index", "Global" is unselected and "1" 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	<input checked="" type="checkbox"/>
Authentication Key
Confirm Authentication Key

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

Figure 155: Configuring Remote Authentication Server (TACACS+)

Security > AAA > Server

Step: 1. Configure Server

Server Type RADIUS TACACS+

Global | Server Index: 1

Server IP Address

Authentication Server TCP Port (1-65535)

Authentication Timeout (1-540) sec

Authentication Retries (1-30)

Set Key

Authentication Key

Confirm Authentication Key

Apply Revert

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

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

Figure 156: Configuring AAA Server Groups

Security > AAA > Server

Step: 2. Configure Group Action: Add

Server Type RADIUS TACACS+

RADIUS Group Name

Sequence At Priority 1

Sequence At Priority 2

Sequence At Priority 3

Sequence At Priority 4

Sequence At Priority 5

Apply Revert

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

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

Figure 157: Showing AAA Server Groups

The screenshot shows the configuration page for AAA Server Groups. At the top, the breadcrumb is 'Security > AAA > Server'. Below it, there are two dropdown menus: 'Step: 2. Configure Group' and 'Action: Show'. Underneath, there are radio buttons for 'Server Type' with 'RADIUS' selected and 'TACACS+' unselected. Below that, it says 'RADIUS Group List Total: 3'. A table follows with three columns: a checkbox, 'Group Name', and 'Member Index'. The table contains three rows: 'radius' with member index '1, 2, 3, 5', 'radius1' with '3, 5, 1', and 'radius2' with '1, 2, 5'. At the bottom right of the table area are 'Delete' and 'Revert' buttons.

<input type="checkbox"/>	Group Name	Member Index
<input type="checkbox"/>	radius	1, 2, 3, 5
<input type="checkbox"/>	radius1	3, 5, 1
<input type="checkbox"/>	radius2	1, 2, 5

Configuring AAA Accounting

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

Command Usage

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

Parameters

These parameters are displayed:

Configure Global

- ◆ **Periodic Update** - Specifies the interval at which the local accounting service updates information for all users on the system to the accounting server. (Range: 0-2147483647 minutes; where 0 means disabled)

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 251](#)). Any other group name refers to a server group configured on the Security > AAA > Server (Configure Group) page.

Configure Service

- ◆ **Accounting Type** – Specifies the service as 802.1X, Command or Exec as described in the preceding section.
 - **802.1X**
 - **Method Name** – Specifies a user defined accounting method to apply to an interface. This method must be defined in the Configure Method page. (Range: 1-255 characters)
 - **Exec**
 - **Console Method Name** – Specifies a user defined method name to apply to console connections.
 - **Method Name** – Specifies a user defined method name to apply to Telnet connections.

Show Information – Summary

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

Show Information – Statistics

- ◆ **User Name** - Displays a registered user name.

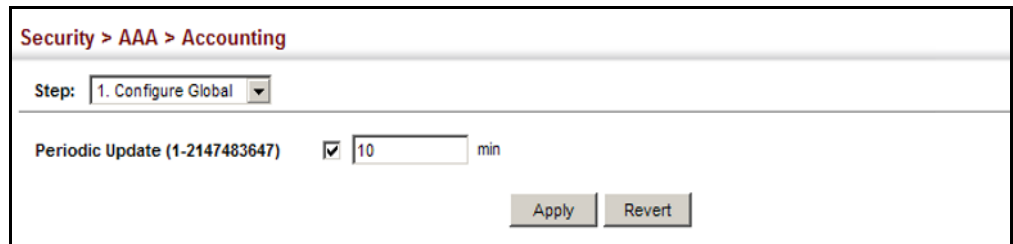
- ◆ **Accounting Type** - Displays the accounting service.
- ◆ **Interface** - Displays the receive port number through which this user accessed the switch.
- ◆ **Time Elapsed** - Displays the length of time this entry has been active.

Web Interface

To configure global settings for AAA accounting:

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

Figure 158: Configuring Global Settings for AAA Accounting



The screenshot shows the Cisco Web Interface for configuring AAA Accounting. The breadcrumb navigation is "Security > AAA > Accounting". The "Step:" dropdown menu is set to "1. Configure Global". Below this, there is a checkbox labeled "Periodic Update (1-2147483647)" which is checked. To the right of the checkbox is a text input field containing the number "10" followed by the unit "min". 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 159: Configuring AAA Accounting Methods

The screenshot shows the configuration page for AAA Accounting. The breadcrumb is "Security > AAA > Accounting". The "Step" is "2. Configure Method" and the "Action" is "Add". The fields are: Accounting Type (802.1X), Method Name (default), Accounting Notice (Start-Stop), and Server Group Name (radius). There are "Apply" and "Revert" buttons at the bottom right.

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 160: Showing AAA Accounting Methods

The screenshot shows the configuration page for AAA Accounting with the "Action" set to "Show". It displays a table of configured accounting methods. The table has columns for Accounting Type, Method Name, Accounting Notice, and Server Group Name. There are "Delete" and "Revert" buttons at the bottom right.

	Accounting Type	Method Name	Accounting Notice	Server Group Name
<input type="checkbox"/>	802.1X	default	Start-Stop	radius
<input type="checkbox"/>	EXEC	default	Start-Stop	tacacs+

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 161: Configuring AAA Accounting Service for 802.1X Service

Security > AAA > Accounting

Step: 3. Configure Service

Accounting Type 802.1X EXEC

Port Method List Total: 28 1 2 3

Port	Method Name
1	default
2	default
3	default
4	
5	

Figure 162: Configuring AAA Accounting Service for Exec Service

Security > AAA > Accounting

Step: 3. Configure Service

Accounting Type 802.1X EXEC

Console Method Name

VTY Method Name

Apply Revert

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

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

Figure 163: Displaying a Summary of Applied AAA Accounting Methods

Security > AAA > Accounting

Step: 4. Show Information

Summary Statistics

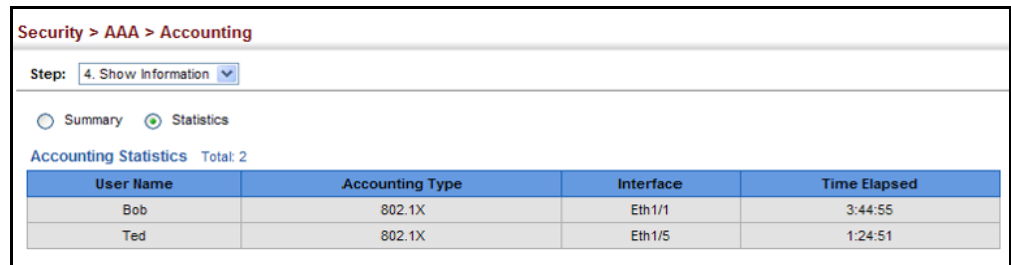
Method List Total: 2

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 164: 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: 1-255 characters)
- ◆ **Server Group Name** - Specifies the authorization server group. (Range: 1-255 characters)

The group name “tacacs+” specifies all configured TACACS+ hosts (see [“Configuring Local/Remote Logon Authentication” on page 251](#)). 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

- ◆ **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 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 165: Configuring AAA Authorization Methods

The screenshot shows a web interface for configuring AAA Authorization. The breadcrumb path is "Security > AAA > Authorization". The "Step" dropdown is set to "1. Configure Method" and the "Action" dropdown is set to "Add". The configuration fields are as follows:

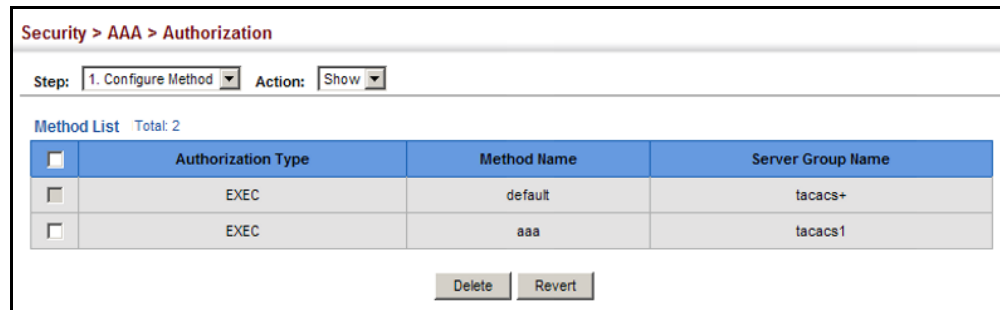
Authorization Type	EXEC
Method Name	default
Server Group Name	<input checked="" type="radio"/> tacacs+ <input type="radio"/> [empty]

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 166: Showing AAA Authorization Methods



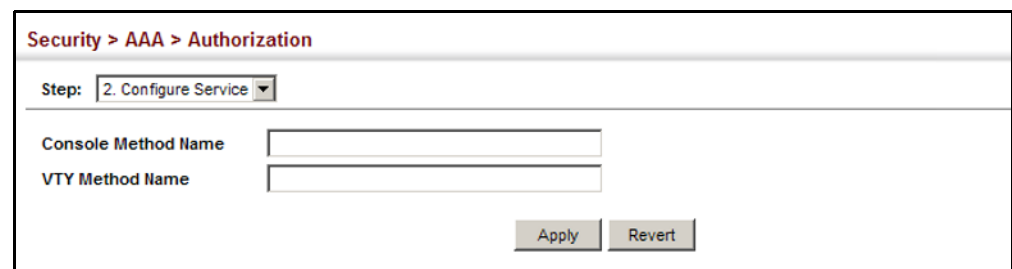
The screenshot shows the 'Security > AAA > Authorization' configuration page. At the top, there is a breadcrumb trail and a navigation bar with 'Step: 1. Configure Method' and 'Action: Show'. Below this is a 'Method List' section with a 'Total: 2' indicator. The table below has three columns: 'Authorization Type', 'Method Name', and 'Server Group Name'. There are two rows of data, each with a checkbox in the first column. At the bottom of the table are 'Delete' and 'Revert' buttons.

<input type="checkbox"/>	Authorization Type	Method Name	Server Group Name
<input type="checkbox"/>	EXEC	default	tacacs+
<input type="checkbox"/>	EXEC	aaa	tacacs1

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 167: Configuring AAA Authorization Methods for Exec Service



The screenshot shows the 'Security > AAA > Authorization' configuration page at the '2. Configure Service' step. It features two input fields: 'Console Method Name' and 'VTY Method Name'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 168: Displaying the Applied AAA Authorization Method

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

Configuring User Accounts

Use the Security > User Accounts page to control management access to the switch based on manually configured user names and passwords.

Command Usage

- ◆ The default guest name is “guest” with the password “guest.” The default administrator name is “admin” with the password “admin.”
- ◆ The guest only has read access for most configuration parameters. However, the administrator has write access for all parameters governing the onboard agent. You should therefore assign a new administrator password as soon as possible, and store it in a safe place.

Parameters

These parameters are displayed:

- ◆ **User Name** – The name of the user.
(Maximum length: 32 characters; maximum number of users: 16)
- ◆ **Access Level** – Specifies the user level. (Options: 0 - Normal, 15 - Privileged)
Normal privilege level provides access to a limited number of the commands which display the current status of the switch, as well as several database clear and reset functions. Privileged level provides full access to all commands.
- ◆ **Password Type** – Plain Text or Encrypted password.
 - **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. There is no need for you to manually configure encrypted passwords.

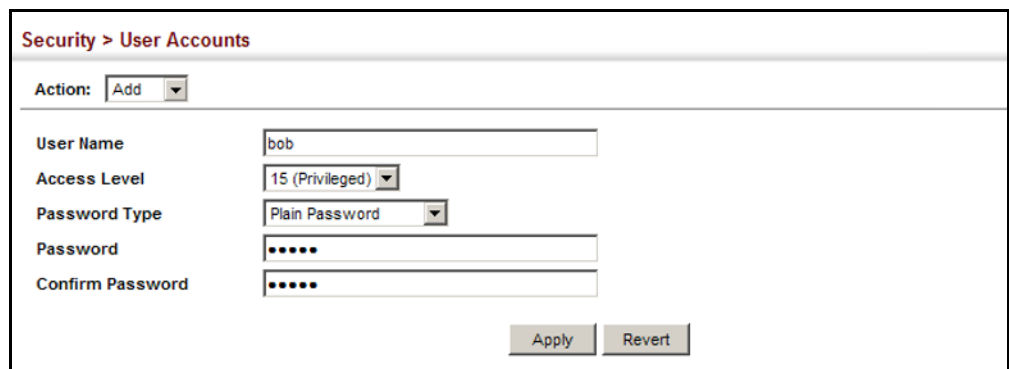
- ◆ **Password** – Specifies the user password. (Range: 0-32 characters, case sensitive)
- ◆ **Confirm Password** – Re-type the string entered in the previous field to ensure no errors were made. The switch will not change the password if these two fields do not match.

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 169: Configuring User Accounts



The screenshot shows a web interface for configuring user accounts. At the top, it says "Security > User Accounts". Below that is an "Action:" dropdown menu with "Add" selected. The 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 five dots representing a masked password.
- Confirm Password:** A text input field with five dots representing a masked password.

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

To show user accounts:

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

Figure 170: Showing User Accounts

The screenshot shows a web interface for managing user accounts. At the top, it says "Security > User Accounts". Below that is an "Action:" dropdown menu set to "Show". The main content is a table titled "User Account List Total: 3". The table has three columns: a checkbox, "User Name", and "Access Level". There are three rows of data: "admin" with access level 15, "guest" with access level 0, and "bob" with access level 15. At the bottom right of the table area are "Delete" and "Revert" buttons.

<input type="checkbox"/>	User Name	Access Level
<input type="checkbox"/>	admin	15
<input type="checkbox"/>	guest	0
<input type="checkbox"/>	bob	15

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 251.)

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, or 0 for disabled; Default: 3600 seconds)

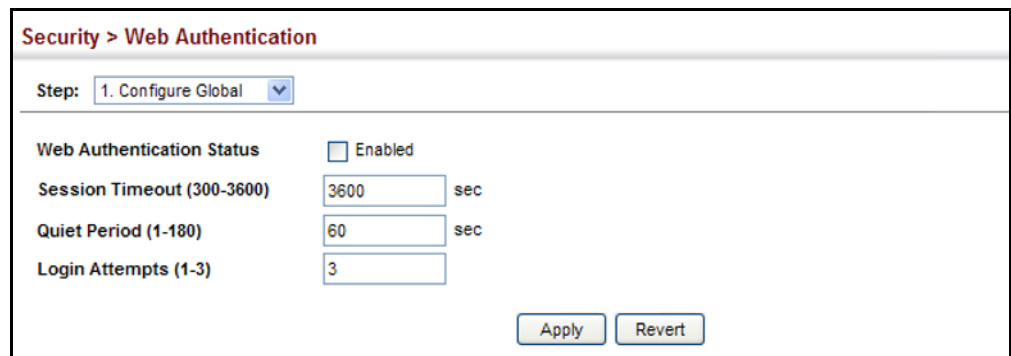
- ◆ **Quiet Period** – Configures how long a host must wait to attempt authentication again after it has exceeded the maximum allowable failed login attempts. (Range: 1-180 seconds; Default: 60 seconds)
- ◆ **Login Attempts** – Configures the amount of times a supplicant may attempt and fail authentication 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 171: Configuring Global Settings for Web Authentication



The screenshot shows the configuration page for Web Authentication. The breadcrumb is "Security > Web Authentication". The "Step" dropdown is set to "1. Configure Global". The "Web Authentication Status" is currently disabled (checkbox is unchecked). 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. There are "Apply" and "Revert" buttons at the bottom right.

Parameter	Value
Web Authentication Status	<input type="checkbox"/> Enabled
Session Timeout (300-3600)	3600 sec
Quiet Period (1-180)	60 sec
Login Attempts (1-3)	3

Configuring Interface Settings for Web Authentication

Use the Security > Web Authentication (Configure Interface) page to enable web authentication on a port, and display information for any connected hosts.

Parameters

These parameters are displayed:

- ◆ **Port** – Indicates the port being configured.
- ◆ **Status** – Configures the web authentication status for the port.
- ◆ **Host IP Address** – Indicates the IP address of each connected host.
- ◆ **Remaining Session Time** – Indicates the remaining time until the current authorization session for the host expires.

- ◆ **Apply** – Enables web authentication if the Status box is checked.
- ◆ **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 172: Configuring Interface Settings for Web Authentication

The screenshot displays the configuration interface for web authentication. At the top, the breadcrumb is 'Security > Web Authentication'. The current step is '2. Configure Interface'. The 'Port' is set to '1' and the 'Status' is 'Enabled'. There are 'Apply' and 'Revert' buttons. Below this is the 'Authenticated Host List' with a total of 2 hosts. The table has columns for 'Host IP Address' and 'Remaining Session Time (sec)'. The first host is 10.1.1.1 with 300 seconds remaining, and the second is 10.2.2.2 with 100 seconds remaining. At the bottom, there are 'Re-authenticate' and 'Revert' buttons.

	Host IP Address	Remaining Session Time (sec)
<input type="checkbox"/>	10.1.1.1	300
<input type="checkbox"/>	10.2.2.2	100

Network Access (MAC Address Authentication)

Some devices connected to switch ports may not be able to support 802.1X authentication due to hardware or software limitations. This is often true for devices such as network printers, IP phones, and some wireless access points. The switch enables network access from these devices to be controlled by authenticating device MAC addresses with a central RADIUS server.



Note: RADIUS authentication must be activated and configured properly for the MAC Address authentication feature to work properly. (See [“Configuring Remote Logon Authentication Servers”](#) on page 252.)

Note: MAC authentication cannot be configured on trunk ports.

Command Usage

- ◆ MAC address authentication controls access to the network by authenticating the MAC address of each host that attempts to connect to a switch port. Traffic received from a specific MAC address is forwarded by the switch only if the source MAC address is successfully authenticated by a central RADIUS server. While authentication for a MAC address is in progress, all traffic is blocked until authentication is completed. On successful authentication, the RADIUS server may optionally assign VLAN and quality of service settings for the switch port.
- ◆ When enabled on a port, the authentication process sends a Password Authentication Protocol (PAP) request to a configured RADIUS server. The user name and password are both equal to the MAC address being authenticated. On the RADIUS server, PAP user name and passwords must be configured in the MAC address format XX-XX-XX-XX-XX-XX (all in upper case).
- ◆ Authenticated MAC addresses are stored as dynamic entries in the switch secure MAC address table and are removed when the aging time expires. The maximum number of secure MAC addresses supported for the switch system is 1024.
- ◆ Configured static MAC addresses are added to the secure address table when seen on a switch port. Static addresses are treated as authenticated without sending a request to a RADIUS server.
- ◆ When port status changes to down, all MAC addresses mapped to that port are cleared from the secure MAC address table. Static VLAN assignments are not restored.
- ◆ The RADIUS server may optionally return a VLAN identifier list to be applied to the switch port. The following attributes need to be configured on the RADIUS server.
 - **Tunnel-Type** = VLAN
 - **Tunnel-Medium-Type** = 802

- **Tunnel-Private-Group-ID** = 1u,2t [VLAN ID list]

The VLAN identifier list is carried in the RADIUS “Tunnel-Private-Group-ID” attribute. The VLAN list can contain multiple VLAN identifiers in the format “1u,2t,3u” where “u” indicates an untagged VLAN and “t” a tagged VLAN.

- ◆ The RADIUS server may optionally return dynamic QoS assignments to be applied to a switch port for an authenticated user. The “Filter-ID” attribute (attribute 11) can be configured on the RADIUS server to pass the following QoS information:

Table 18: Dynamic QoS Profiles

Profile	Attribute Syntax	Example
DiffServ	service-policy-in = <i>policy-map-name</i>	service-policy-in=p1
Rate Limit	rate-limit-input = <i>rate</i>	rate-limit-input=100 (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 326](#)).

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 a connected host must be reauthenticated. When the reauthentication time expires for a secure MAC address, it is reauthenticated with the RADIUS server. During the reauthentication process traffic through the port remains unaffected. (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 173: Configuring Global Settings for Network Access

The screenshot shows a web interface for configuring network access. The breadcrumb is 'Security > Network Access'. The 'Step' dropdown is set to '1. Configure Global'. Under 'Aging Status', there is a checked checkbox and the text 'Enabled'. Below that, 'Reauthentication Time (120-1000000)' is set to '30000' with 'sec' as the unit. At the bottom right, there are 'Apply' and 'Revert' buttons.

Configuring Network Access for Ports

Use the Security > Network Access (Configure Interface - General) page to configure MAC authentication on switch ports, including enabling address authentication, setting the maximum MAC count, and enabling dynamic VLAN or dynamic QoS assignments.

Parameters

These parameters are displayed:

◆ MAC Authentication

- **Status** – Enables MAC authentication on a port. (Default: Disabled)
- **Intrusion** – Sets the port response to a host MAC authentication failure to either block access to the port or to pass traffic through. (Options: Block, Pass; Default: Block)
- **Max MAC Count**⁵ – 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**⁵ – Sets the maximum number of MAC addresses that can be authenticated on a port interface via all forms of

5. 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.

authentication (including Network Access and IEEE 802.1X). (Range: 1-1024; Default: 1024)

- ◆ **Guest VLAN** – Specifies the VLAN to be assigned to the port when 802.1X Authentication fails. (Range: 0-4094, where 0 means disabled; Default: Disabled)

The VLAN must already be created and active (see [“Configuring VLAN Groups” on page 138](#)). 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 326](#)).

- ◆ **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 174: Configuring Interface Settings for Network Access

The screenshot shows the 'Security > Network Access' configuration page. The 'Step: 2. Configure Interface' dropdown is selected. The 'General' tab is active. Below the tabs is a 'Port List' with a total of 28 ports. A table displays the configuration for the first five ports. The table has columns for Port, Status, Intrusion, Max MAC Count (1-1024), Network Access Max MAC Count (1-2048), Guest VLAN (0-4094, 0: Disabled), Dynamic VLAN, Dynamic QoS, and MAC Filter ID (1-64). All ports are shown with 'Enabled' status and 'Block' intrusion. Max MAC counts are set to 1024. Network Access Max MAC counts are 1024. Guest VLAN is 0. Dynamic VLAN is 'Enabled'. Dynamic QoS is 'Enabled'. MAC Filter ID is empty.

Port	MAC Authentication			Network Access Max MAC Count (1-2048)	Guest VLAN (0-4094, 0: Disabled)	Dynamic VLAN	Dynamic QoS	MAC Filter ID (1-64)
	Status	Intrusion	Max MAC Count (1-1024)					
1	<input type="checkbox"/> Enabled	Block	1024	1024	0	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	
2	<input checked="" type="checkbox"/> Enabled	Block	1024	1024	0	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	
3	<input type="checkbox"/> Enabled	Block	1024	1024	0	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	
4	<input type="checkbox"/> Enabled	Block	1024	1024	0	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	
5	<input type="checkbox"/> Enabled	Block	1024	1024	0	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/> 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.

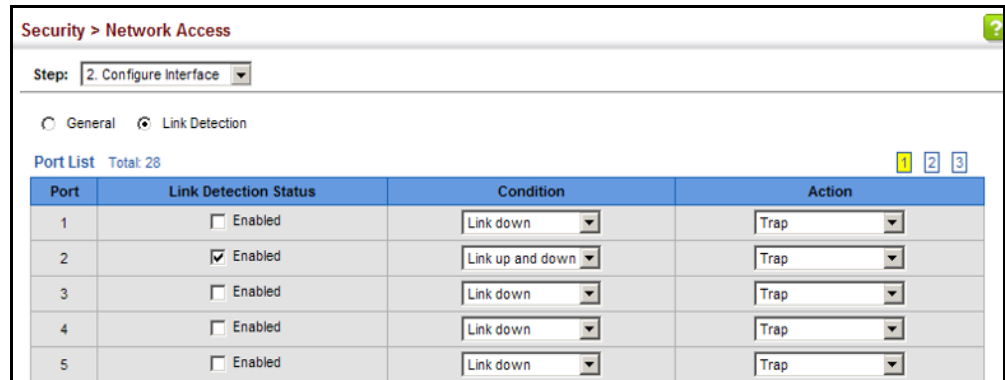
Web Interface

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 175: Configuring Link Detection for Network Access



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.
- ◆ **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 176: 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 177: Showing the MAC Address Filter Table for Network Access

Security > Network Access

Step: 3. Configure MAC Filter Action: Show

MAC Filter List Total: 1

<input type="checkbox"/>	Filter ID	MAC Address	MAC Address Mask
<input type="checkbox"/>	22	11-22-33-44-55-66	FF-FF-FF-FF-FF-FF

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 178: Showing Addresses Authenticated for Network Access

The screenshot shows the 'Security > Network Access' configuration page. At the top, the 'Step' is set to '4. Show Information'. Under 'Query by:', the 'Sort Key' is set to 'MAC Address'. There are three checkboxes: 'MAC Address', 'Interface', and 'Attribute', all of which are unchecked. Below these are input fields for 'MAC Address', 'Interface' (set to '1'), and 'Attribute' (set to 'Static'). A 'Query' button is located to the right. Below the query options, the text 'Authenticated MAC Address List Total: 8' is displayed. A table with 6 columns (checkbox, MAC Address, Interface, RADIUS Server, Time, Attribute) lists 5 entries of authenticated MAC addresses.

<input type="checkbox"/>	MAC Address	Interface	RADIUS Server	Time	Attribute
<input type="checkbox"/>	00-00-86-45-F2-23	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 16m 12s	Dynamic
<input type="checkbox"/>	00-00-E8-5E-E1-DD	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 32m 24s	Dynamic
<input type="checkbox"/>	00-00-E8-81-93-30	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 40m 32s	Dynamic
<input type="checkbox"/>	00-01-80-31-B8-30	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 18m 51s	Dynamic
<input type="checkbox"/>	00-01-80-36-95-D8	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 32m 22s	Dynamic

Configuring HTTPS

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

Configuring Global Settings for HTTPS

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

Command Usage

- ◆ HTTP and HTTPS are implemented as mutually exclusive services on the switch. (HTTP can only be configured through the CLI using the "ip http server" command described in the *CLI Reference Guide*.)
- ◆ If you enable HTTPS, you must indicate this in the URL that you specify in your browser: `https://device[:port_number]`
- ◆ When you start HTTPS, the connection is established in this way:
 - The client authenticates the server using the server's digital certificate.
 - The client and server negotiate a set of security protocols to use for the connection.
 - The client and server generate session keys for encrypting and decrypting data.
- ◆ The client and server establish a secure encrypted connection.
A padlock icon should appear in the status bar for Internet Explorer 6.x or above, or Mozilla Firefox 4.x or above.

- ◆ The following web browsers and operating systems currently support HTTPS:

Table 19: HTTPS System Support

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

- ◆ To specify a secure-site certificate, see [“Replacing the Default Secure-site Certificate” on page 281](#).



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

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

Parameters

These parameters are displayed:

- ◆ **HTTPS Status** – Allows you to enable/disable the HTTPS server feature on the switch. (Default: Disabled)
- ◆ **HTTPS Port** – Specifies the TCP port number used for HTTPS connection to the switch’s web interface. (Default: Port 443)

Web Interface

To configure HTTPS:

1. Click Security, HTTPS.
2. Select Configure Global from the Step list.
3. Enable HTTPS and specify the port number if required.
4. Click Apply.

Figure 179: Configuring HTTPS

The screenshot shows a web interface for configuring HTTPS. At the top, it says "Security > HTTPS". Below that, there is a dropdown menu for "Action" set to "Configure Global". Underneath, there are two settings: "HTTPS Status" which is checked and labeled "Enabled", and "UDP Port (1-65535)" which is set to "443". At the bottom right, there are two buttons: "Apply" and "Revert".

Replacing the Default Secure-site Certificate

Use the Security > HTTPS (Copy Certificate) page to replace the default secure-site certificate.

When you log onto the web interface using HTTPS (for secure access), a Secure Sockets Layer (SSL) certificate appears for the switch. By default, the certificate that the web browser displays will be associated with a warning that the site is not recognized as a secure site. This is because the certificate has not been signed by an approved certification authority. If you want this warning to be replaced by a message confirming that the connection to the switch is secure, you must obtain a unique certificate and a private key and password from a recognized certification authority.



Caution: For maximum security, we recommend you obtain a unique Secure Sockets Layer certificate at the earliest opportunity. This is because the default certificate for the switch is not unique to the hardware you have purchased.

When you have obtained these, place them on your TFTP server and transfer them to the switch to replace the default (unrecognized) certificate with an authorized one.



Note: The switch must be reset for the new certificate to be activated. To reset the switch, see [“Resetting the System” on page 91](#) 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 180: 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, there are five input fields: "TFTP Server IP Address" with the value "192.168.0.4", "Certificate Source File Name" with "ES3510MA-site-certificate", "Private Key Source File Name" with "ES3510MA-private-key", "Private Password" with masked characters, and "Confirm Password" with masked characters. At the bottom right, there 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 251). 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 288](#), 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 265](#).) 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 286](#).

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 181: Configuring the SSH Server

The screenshot shows the 'Security > SSH' configuration page. At the top, there is a breadcrumb 'Security > SSH' and a 'Step:' dropdown menu set to '1. Configure Global'. Below this, the configuration parameters are listed:

SSH Server Status	<input checked="" type="checkbox"/> Enabled
Version	2.0
Authentication Timeout (1-120)	<input type="text" value="120"/> sec
Authentication Retries (1-5)	<input type="text" value="3"/>
Server-Key Size (512-896)	<input type="text" value="768"/>

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

Generating the Host Key Pair

Use the Security > SSH (Configure Host Key - Generate) page to generate a host public/private key pair used to provide secure communications between an SSH client and the switch. After generating this key pair, you must provide the host public key to SSH clients and import the client's public key to the switch as described in the section "Importing User Public Keys" on page 288.



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 285.

Parameters

These parameters are displayed:

- ◆ **Host-Key Type** – The key type used to generate the host key pair (i.e., public and private keys). (Range: RSA (Version 1), DSA (Version 2), Both; Default: Both)
The SSH server uses RSA or DSA for key exchange when the client first establishes a connection with the switch, and then negotiates with the client to select either DES (56-bit) or 3DES (168-bit) for data encryption.



Note: The switch uses only RSA Version 1 for SSHv1.5 clients and DSA Version 2 for SSHv2 clients.

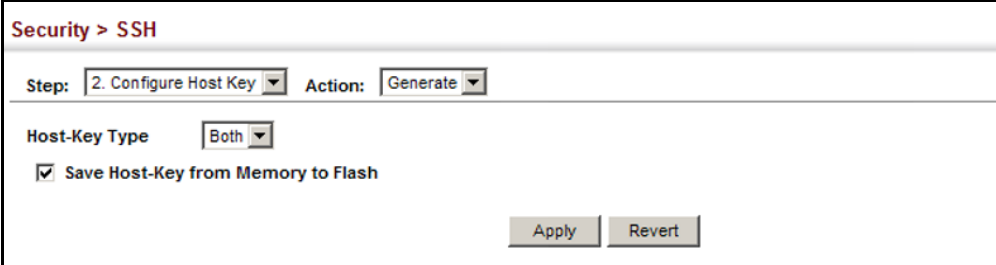
- ◆ **Save Host-Key from Memory to Flash** – Saves the host key from RAM (i.e., volatile memory) to flash memory. Otherwise, the host key pair is stored to RAM by default. Note that you must select this item prior to generating the host-key pair. (Default: Disabled)

Web Interface

To generate the SSH host key pair:

1. Click Security, SSH.
2. Select Configure Host Key from the Step list.
3. Select Generate from the Action list.
4. Select the host-key type from the drop-down box.
5. Select the option to save the host key from memory to flash if required.
6. Click Apply.

Figure 182: Generating the SSH Host Key Pair

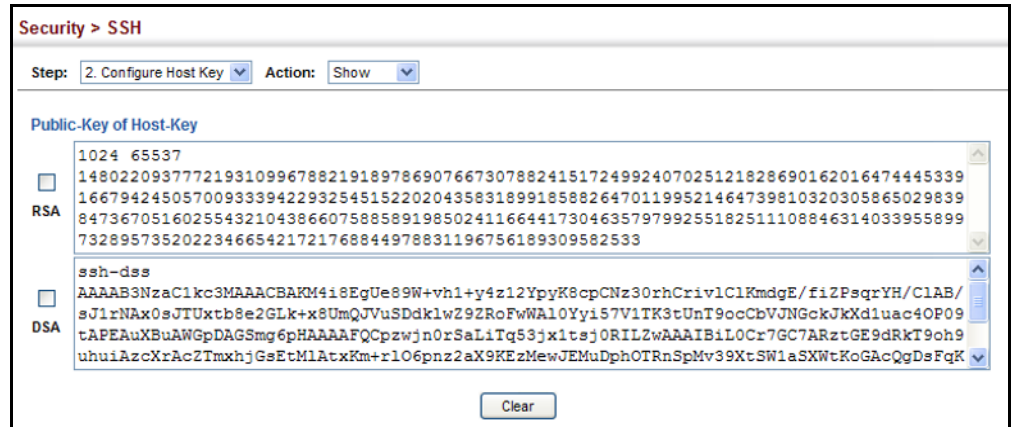


The screenshot shows a web interface for configuring SSH. The breadcrumb is "Security > SSH". The "Step:" dropdown is set to "2. Configure Host Key" and the "Action:" dropdown is set to "Generate". Below this, the "Host-Key Type" dropdown is set to "Both". There is a checked checkbox labeled "Save Host-Key from Memory to Flash". At the bottom right, there are "Apply" and "Revert" buttons.

To display or clear the SSH host key pair:

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

Figure 183: Showing the SSH Host Key Pair



Importing User Public Keys

Use the Security > SSH (Configure User Key - Copy) page to upload a user's public key to the switch. This public key must be stored on the switch for the user to be able to log in using the public key authentication mechanism. If the user's public key does not exist on the switch, SSH will revert to the interactive password authentication mechanism to complete authentication.

Parameters

These parameters are displayed:

- ◆ **User Name** – This drop-down box selects the user who's public key you wish to manage. Note that you must first create users on the User Accounts page (see "Configuring User Accounts" on page 265).
- ◆ **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 184: Copying the SSH User's Public Key

Security > SSH

Step: 3. Configure User Key Action: Copy

User Name: steve

User-Key Type: RSA

TFTP Server IP Address: 192.168.0.61

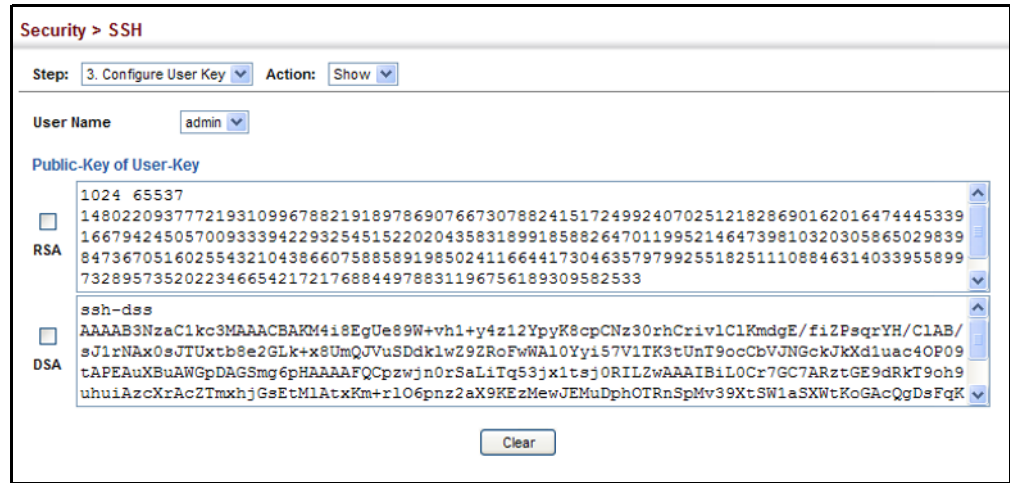
Source File Name: rsa.pub

Apply 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 185: Showing the SSH User's Public Key



Access Control Lists

Access Control Lists (ACL) provide packet filtering for IPv4 frames (based on address, protocol, Layer 4 protocol port number or TCP control code), IPv6 frames (based on address, DSCP traffic class, next header type, or any frames (based on MAC address or Ethernet type). To filter incoming packets, first create an access list, add the required rules, and then bind the list to a specific port.

Configuring Access Control Lists –

An ACL is a sequential list of permit or deny conditions that apply to IP addresses, MAC addresses, or other more specific criteria. This switch tests ingress 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 ACL rules per system is 384 rules.
- ◆ An ACL can have up to 64 rules for IPv4 ACLs, IPv6 ACLs, and ARP ACLs, and up to 45 for MAC ACLs. However, due to resource restrictions, the average number of rules bound to the ports should not exceed 20.
- ◆ 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 efficiency. 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.

The order in which active ACLs are checked is as follows:

1. User-defined rules in IP and MAC ACLs for ingress or egress ports are checked in parallel.
2. Rules within an ACL are checked in the configured order, from top to bottom.
3. If the result of checking an IP ACL is to permit a packet, but the result of a MAC ACL on the same packet is to deny it, the packet will be denied (because the decision to deny a packet has a higher priority for security reasons). A packet will also be denied if the IP ACL denies it and the MAC ACL accepts it.

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

Command Usage

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

Parameters

These parameters are displayed:

Add

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

Add Rule

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

◆ **Mode**

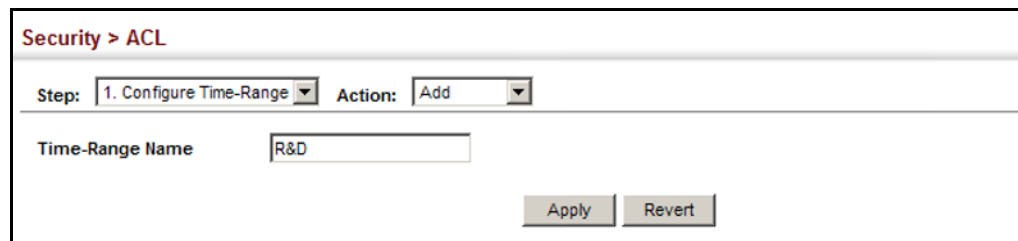
- **Absolute** – Specifies a specific time or time range.
 - **Start/End** – Specifies the hours, minutes, month, day, and year at which to start or end.
- **Periodic** – Specifies a periodic interval.
 - **Start/To** – Specifies the days of the week, hours, and minutes at which to start or end.

Web Interface

To configure a time range:

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

Figure 186: Setting the Name of a Time Range

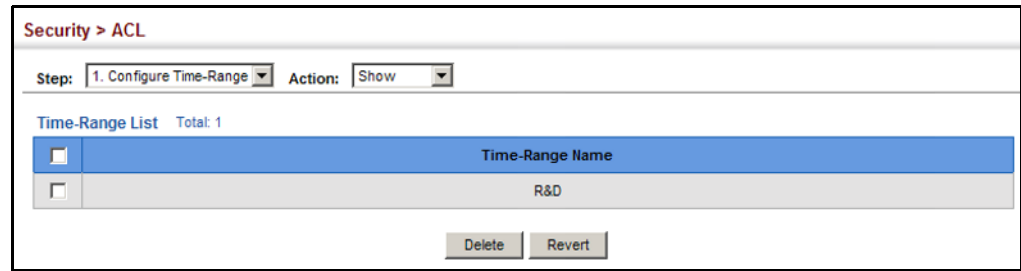


The screenshot shows a web interface for configuring a time range. At the top, it says "Security > ACL". Below that, there are two dropdown menus: "Step:" with "1. Configure Time-Range" selected, and "Action:" with "Add" selected. Below these is a text input field labeled "Time-Range Name" containing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

To show a list of time ranges:

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

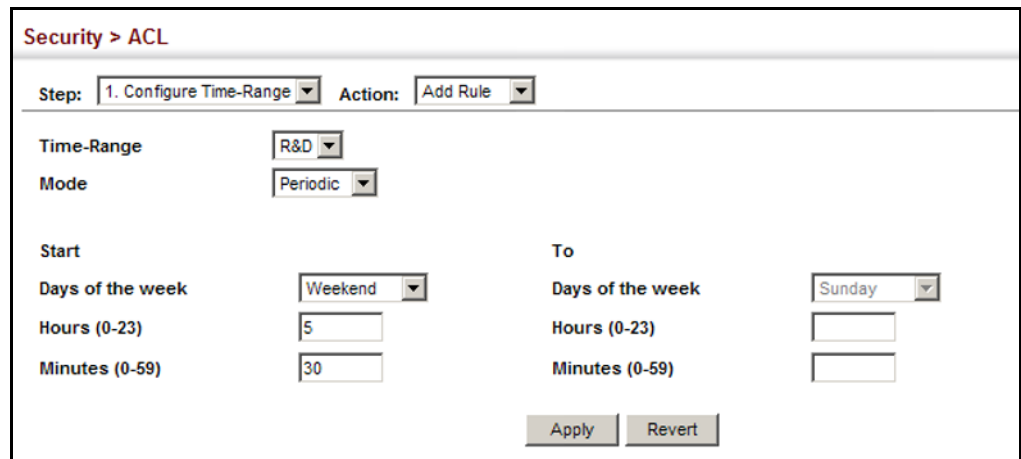
Figure 187: Showing a List of Time Ranges



To configure a rule for a time range:

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

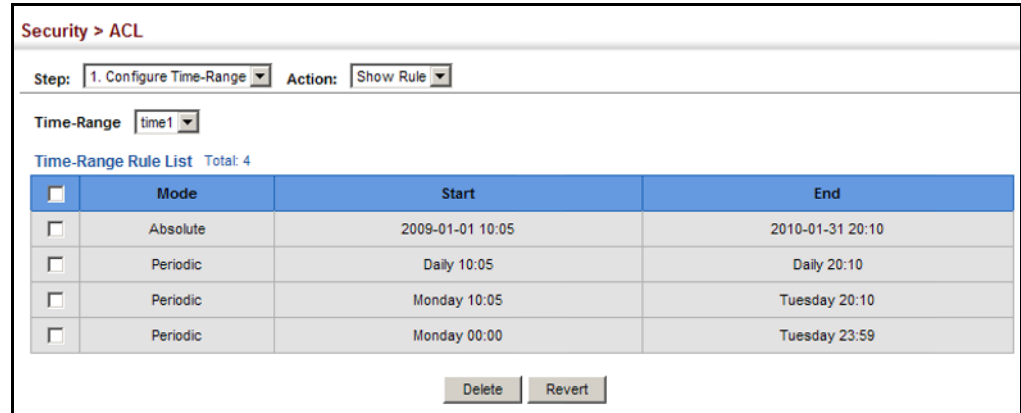
Figure 188: Add a Rule to a Time Range



To show the rules configured for a time range:

1. Click Security, ACL.
2. Select Configure Time Range from the Step list.
3. Select Show Rule from the Action list.

Figure 189: Showing the Rules Configured for a Time Range



Showing TCAM Utilization

Use the Security > ACL (Configure ACL - Show TCAM) page to show utilization parameters for TCAM (Ternary Content Addressable Memory), including the number policy control entries in use, the number of free entries, and the overall percentage of TCAM in use.

Command Usage

Policy control entries (PCEs) are used by various system functions which rely on rule-based searches, including Access Control Lists (ACLs), IP Source Guard filter rules, Quality of Service (QoS) processes, QinQ, MAC-based VLANs, VLAN translation, or traps.

For example, when binding an ACL to a port, each rule in an ACL will use two PCEs; and when setting an IP Source Guard filter rule for a port, the system will also use two PCEs.

Parameters

These parameters are displayed:

- ◆ **Total Policy Control Entries** – The number policy control entries in use.
- ◆ **Free Policy Control Entries** – The number of policy control entries available for use.
- ◆ **Entries Used by System** – The number of policy control entries used by the operating system.
- ◆ **Entries Used by User** – The number of policy control entries used by configuration settings, such as access control lists.
- ◆ **TCAM Utilization** – The overall percentage of TCAM in use.

Web Interface

To show information on TCAM utilization:

1. Click Security, ACL.
2. Select Configure ACL from the Step list.
3. Select Show TCAM from the Action list.

Figure 190: Showing TCAM Utilization

Security > ACL	
Step:	2. Configure ACL
Action:	Show TCAM
Total Policy Control Entries	384
Free Policy Control Entries	290
Entries Used by System	94
Entries Used by User	0
TCAM Utilization	24.47%

Setting the ACL Name and Type

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

Parameters

These parameters are displayed:

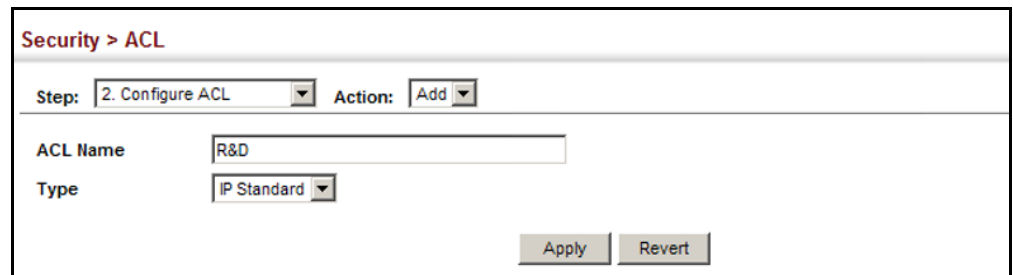
- ◆ **ACL Name** – Name of the ACL. (Maximum length: 32 characters)
- ◆ **Type** – The following filter modes are supported:
 - **IP Standard:** IPv4 ACL mode filters packets based on the source IPv4 address.
 - **IP Extended:** IPv4 ACL mode filters packets based on the source or destination IPv4 address, 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 312](#)).

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 191: Creating an ACL

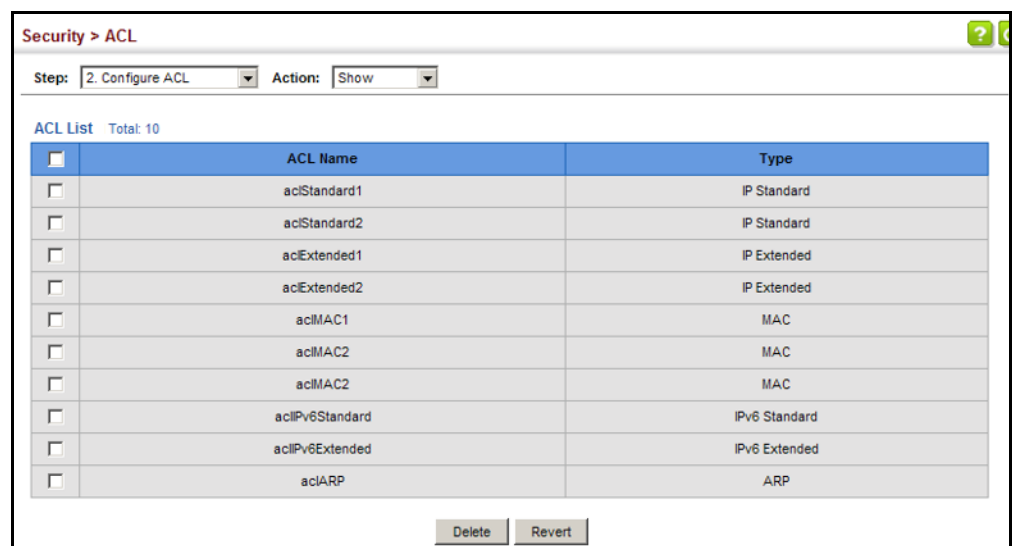


The screenshot shows the 'Security > ACL' configuration page. At the top, there is a breadcrumb 'Security > ACL'. Below it, there are two dropdown menus: 'Step: 2. Configure ACL' and 'Action: Add'. The main form has two fields: 'ACL Name' with the value 'R&D' and 'Type' with the value 'IP Standard'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

To show a list of ACLs:

1. Click Security, ACL.
2. Select Configure ACL from the Step list.
3. Select Show from the Action list.

Figure 192: Showing a List of ACLs



The screenshot shows the 'Security > ACL' configuration page with the 'Action' dropdown set to 'Show'. Below the form, there is a table titled 'ACL List' with a total of 10 items. The table has three columns: 'ACL Name', 'Type', and a checkbox column. The table lists various ACLs and their types.

<input type="checkbox"/>	ACL Name	Type
<input type="checkbox"/>	aclStandard1	IP Standard
<input type="checkbox"/>	aclStandard2	IP Standard
<input type="checkbox"/>	aclExtended1	IP Extended
<input type="checkbox"/>	aclExtended2	IP Extended
<input type="checkbox"/>	aclMAC1	MAC
<input type="checkbox"/>	aclMAC2	MAC
<input type="checkbox"/>	aclMAC2	MAC
<input type="checkbox"/>	aclIPv6Standard	IPv6 Standard
<input type="checkbox"/>	aclIPv6Extended	IPv6 Extended
<input type="checkbox"/>	aclARP	ARP

At the bottom right of the table, there are two buttons: 'Delete' and '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 rules which permit or deny a packet, or re-direct a packet to another port.
- ◆ **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 193: Configuring a Standard IPv4 ACL

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

Configuring an Extended IPv4 ACL Use the Security > ACL (Configure ACL - Add Rule - IP Extended) page to configure an Extended IPv4 ACL.

Parameters

These parameters are displayed:

- ◆ **Type** – Selects the type of ACLs to show in the Name list.
- ◆ **Name** – Shows the names of ACLs matching the selected type.
- ◆ **Action** – An ACL can contain any combination of permit or deny rules.
- ◆ **Source/Destination Address Type** – Specifies the source or destination IP address type. Use “Any” to include all possible addresses, “Host” to specify a specific host address in the Address field, or “IP” to specify a range of addresses with the Address and Subnet Mask fields. (Options: Any, Host, IP; Default: Any)
- ◆ **Source/Destination IP Address** – Source or destination IP address.
- ◆ **Source/Destination Subnet Mask** – Subnet mask for source or destination address. (See the description for Subnet Mask on [page 297](#).)
- ◆ **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: TCP)

- ◆ **Service Type** – Packet priority settings based on the following criteria:
 - **ToS** – Type of Service level. (Range: 0-15)
 - **Precedence** – IP precedence level. (Range: 0-7)
 - **DSCP** – DSCP priority level. (Range: 0-63)
- ◆ **Control Code** – Decimal number (representing a bit string) that specifies flag bits in byte 14 of the TCP header. (Range: 0-63)
- ◆ **Control Code Bit Mask** – Decimal number representing the code bits to match. (Range: 0-63)

The control bit mask is a decimal number (for an equivalent binary bit mask) that is applied to the control code. Enter a decimal number, where the equivalent binary bit “1” means to match a bit and “0” means to ignore a bit. The following bits may be specified:

- 1 (fin) – Finish
- 2 (syn) – Synchronize
- 4 (rst) – Reset
- 8 (psh) – Push
- 16 (ack) – Acknowledgement
- 32 (urg) – Urgent pointer

For example, use the code value and mask below to catch packets with the following flags set:

- SYN flag valid, use control-code 2, control bit mask 2
- Both SYN and ACK valid, use control-code 18, control bit mask 18
- SYN valid and ACK invalid, use control-code 2, control bit mask 18

- ◆ **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 194: Configuring an Extended IPv4 ACL

The screenshot shows the 'Security > ACL' configuration page. The 'Step' is '2. Configure ACL' and the 'Action' is 'Add Rule'. The 'Type' is 'IP Extended'. The 'Name' is 'R&D#2'. The 'Action' is 'Permit'. The 'Source Address Type' is 'IP', with 'Source IP Address' set to '10.7.1.0' and 'Source Subnet Mask' set to '255.255.255.0'. The 'Destination Address Type' is 'Any', with 'Destination IP Address' and 'Destination Subnet Mask' both set to '0.0.0.0'. The 'Protocol' is 'TCP (6)'. The 'Service Type' is 'Precedence (0-7)'. There are 'Apply' and 'Revert' buttons at the bottom right.

Configuring a Standard IPv6 ACL Use the Security > ACL (Configure ACL - Add Rule - IPv6 Standard) page to configure a Standard IPv6 ACL.

Parameters

These parameters are displayed:

- ◆ **Type** – Selects the type of ACLs to show in the Name list.
- ◆ **Name** – Shows the names of ACLs matching the selected type.
- ◆ **Action** – An ACL can contain any combination of permit or deny rules.
- ◆ **Source Address Type** – Specifies the source IP address. Use "Any" to include all possible addresses, "Host" to specify a specific host address in the Address field, or "IPv6-Prefix" to specify a range of addresses. (Options: Any, Host, IPv6-Prefix; Default: Any)

- ◆ **Source IPv6 Address** – An IPv6 source address or network class. The address must be formatted according to RFC 2373 “IPv6 Addressing Architecture,” using 8 colon-separated 16-bit hexadecimal values. One double colon may be used in the address to indicate the appropriate number of zeros required to fill the undefined fields.
- ◆ **Source Prefix-Length** – A decimal value indicating how many contiguous bits (from the left) of the address comprise the prefix (i.e., the network portion of the address). (Range: 0-128 bits)
- ◆ **Time Range** – Name of a time range.

Web Interface

To add rules to a Standard IPv6 ACL:

1. Click Security, ACL.
2. Select Configure ACL from the Step list.
3. Select Add Rule from the Action list.
4. Select IPv6 Standard from the Type list.
5. Select the name of an ACL from the Name list.
6. Specify the action (i.e., Permit or Deny).
7. Select the source address type (Any, Host, or IPv6-prefix).
8. If you select “Host,” enter a specific address. If you select “IPv6-prefix,” enter a subnet address and the prefix length.
9. Click Apply.

Figure 195: Configuring a Standard IPv6 ACL

The screenshot shows the 'Security > ACL' configuration page. At the top, it indicates 'Step: 2. Configure ACL' and 'Action: Add Rule'. Below this, there are radio buttons for 'Type': IP Standard, IP Extended, MAC, IPv6 Standard (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 set to 'R&D'. 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.
- ◆ **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 “IPv6-Prefix” to specify a range of addresses. (Options: Any, Host, IPv6-Prefix; Default: Any)
- ◆ **Source/Destination IPv6 Address** – An IPv6 address or network class. The address must be formatted according to RFC 2373 “IPv6 Addressing Architecture,” using 8 colon-separated 16-bit hexadecimal values. One double colon may be used in the address to indicate the appropriate number of zeros required to fill the undefined fields.
- ◆ **Source/Destination Prefix-Length** – A decimal value indicating how many contiguous bits (from the left) of the address comprise the prefix; i.e., the network portion of the address. (Range: 0-128 bits)
- ◆ **DSCP** – DSCP traffic class. (Range: 0-63)

- ◆ **Next Header** – Identifies the type of header immediately following the IPv6 header. (Range: 0-255)

Optional internet-layer information is encoded in separate headers that may be placed between the IPv6 header and the upper-layer header in a packet. There are a small number of such extension headers, each identified by a distinct Next Header value. IPv6 supports the values defined for the IPv4 Protocol field in RFC 1700, and includes these commonly used headers:

- 0: Hop-by-Hop Options (RFC 2460)
- 6: TCP Upper-layer Header (RFC 1700)
- 17: UDP Upper-layer Header (RFC 1700)
- 43: Routing (RFC 2460)
- 44: Fragment (RFC 2460)
- 50: Encapsulating Security Payload (RFC 2406)
- 51: Authentication (RFC 2402)
- 60: Destination Options (RFC 2460)

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

Web Interface

To add rules to an Extended IPv6 ACL:

1. Click Security, ACL.
2. Select Configure ACL from the Step list.
3. Select Add Rule from the Action list.
4. Select IPv6 Extended from the Type list.
5. Select the name of an ACL from the Name list.
6. Specify the action (i.e., Permit or Deny).
7. Select the address type (Any, 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. Set any other required criteria, such as DSCP or next header type.
10. Click Apply.

Figure 196: Configuring an Extended IPv6 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 'IPv6 Extended'. The 'Name' is 'ipv6e'. The 'Action' is 'Permit'. The 'Source Address Type' is 'Host', with a 'Source IPv6 Address' of '2009:db9:2229::79' and a 'Source Prefix Length (0-128)' of '128'. The 'Destination Address Type' is 'IPv6-Prefix', with a 'Destination IPv6 Address' of '2009:db9:2229::78' and a 'Destination Prefix Length (0-128)' of '8'. The 'DSCP (0-63)' and 'Next-Header (0-255)' fields are empty. The 'Time-Range' checkbox is unchecked, and the 'Time-Range' dropdown is set to 'rd'. 'Apply' and 'Revert' buttons are at the bottom right.

Configuring a MAC ACL Use the Security > ACL (Configure ACL - Add Rule - MAC) page to configure a MAC ACL based on hardware addresses, packet format, and Ethernet type.

Parameters

These parameters are displayed:

- ◆ **Type** – Selects the type of ACLs to show in the Name list.
- ◆ **Name** – Shows the names of ACLs matching the selected type.
- ◆ **Action** – An ACL can contain any combination of permit or deny rules.
- ◆ **Source/Destination Address Type** – Use “Any” to include all possible addresses, “Host” to indicate a specific MAC address, or “MAC” to specify an address range with the Address and Bit Mask fields. (Options: Any, Host, MAC; Default: Any)
- ◆ **Source/Destination MAC Address** – Source or destination MAC address.
- ◆ **Source/Destination Bit Mask** – Hexadecimal mask for source or destination MAC address.
- ◆ **VLAN Header** – Specifies the packet header as untagged, tagged, or either.
- ◆ **Packet Format** – This attribute includes the following packet types:
 - **Any** – Any packet format type.

- **Eth2** – Ethernet II packets.
- **LLC-other** – LLC and other packets.
- **SNAP** – SNAP packets.
- ◆ **VID** – VLAN ID. (Range: 1-4094)
- ◆ **VID Bit Mask** – VLAN bit mask. (Range: 0-4095)
- ◆ **Ethernet Type** – This option can only be used to filter Ethernet II formatted packets. (Range: 600-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: 600-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 197: Configuring a MAC ACL

The screenshot shows the 'Security > ACL' configuration page. At the top, it indicates 'Step: 2. Configure ACL' and 'Action: Add Rule'. The 'Type' section has radio buttons for 'IP Standard', 'IP Extended', 'MAC' (selected), 'IPv6 Standard', 'IPv6 Extended', and 'ARP'. The 'Name' field contains 'R&D#3'. The 'Action' dropdown is set to 'Permit'. The 'Source Address Type' is 'Any', and the 'Destination Address Type' is 'Any'. Both 'Source MAC Address' and 'Destination MAC Address' fields contain '00-00-00-00-00-00'. Both 'Source Bit Mask' and 'Destination Bit Mask' fields also contain '00-00-00-00-00-00'. The 'VLAN Header' is set to 'Any'. The 'Packet Format' is 'Any'. There are empty fields for 'VID (1-4094)', 'Ethernet Type', and 'Ethernet Type Bit Mask'. A 'Time-Range' checkbox is unchecked, and the 'rd' dropdown is visible. 'Apply' and 'Revert' buttons are at the bottom right.

Configuring an ARP ACL Use the Security > ACL (Configure ACL - Add Rule - ARP) page to configure ACLs based on ARP message addresses. ARP Inspection can then use these ACLs to filter suspicious traffic (see [“Configuring Global Settings for ARP Inspection” on page 313](#)).

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 297](#).)

- ◆ **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 198: Configuring an ARP ACL

The screenshot shows the 'Security > ACL' configuration page. At the top, the 'Step' is '2. Configure ACL' and the 'Action' is 'Add Rule'. The 'Type' is set to 'ARP'. The 'Name' is 'R&D#7ARP'. The 'Action' is 'Permit'. The 'Packet Type' is '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 which is unchecked. At the bottom right, there are 'Apply' and 'Revert' buttons.

Binding a Port to an Access Control List

After configuring ACLs, use the Security > ACL (Configure Interface) page to bind the ports that need to filter traffic to the appropriate ACLs. You can assign one IP access list and one MAC access list to any port.

Parameters

These parameters are displayed:

- ◆ **Type** – Selects the type of ACLs to bind to a port.
- ◆ **Port** – Fixed port or SFP module. (Range: 1-12/28)
- ◆ **ACL** – ACL used for ingress packets.
- ◆ **Time Range** – Name of a time range.
- ◆ **Counter** – Enables counter for ACL statistics.

Web Interface

To bind an ACL to a port:

1. Click Security, ACL.
2. Select Configure Interface from the Step list.
3. Select IP or MAC from the Type list.
4. Select a port.
5. Select the name of an ACL from the ACL list.

6. Click Apply.

Figure 199: Binding a Port to an ACL

The screenshot shows the 'Security > ACL' configuration window. At the top, it indicates 'Step: 3. Configure Interface' and 'Action: Configure'. Below this, there are radio buttons for 'Type' (IP, MAC, IPv6), with 'IP' selected. A 'Port' dropdown menu is set to '1'. The configuration is divided into 'Ingress' and 'Egress' sections. In the 'Ingress' section, 'ACL' is checked and set to 'R&D', 'Time-Range' is checked and set to 'R&D', and 'Counter' is checked. In the 'Egress' section, 'ACL', 'Time-Range', and 'Counter' are all unchecked. At the bottom right, there are 'Apply' and 'Revert' buttons.

Configuring ACL Mirroring

After configuring ACLs, use the Security > ACL (Configure Interface – Add Mirror) page to mirror traffic matching an ACL from one or more source ports to a target port for real-time analysis. You can then attach a logic analyzer or RMON probe to the target port and study the traffic crossing the source VLAN(s) in a completely unobtrusive manner.

Command Usage

ACL-based mirroring is only used for ingress traffic. To mirror an ACL, follow these steps:

1. Create an ACL as described in the preceding sections.
2. Add one or more mirrored ports to ACL as described under “[Binding a Port to an Access Control List](#)” on page 308.
3. Use the Add Mirror page to specify the ACL and the destination port to which matching traffic will be mirrored.

Parameters

These parameters are displayed:

- ◆ **Port** – Port identifier.
- ◆ **ACL** – ACL used for ingress packets.

Web Interface

To bind an ACL to a port:

1. Click Security, ACL.
2. Select Configure Interface from the Step list.
3. Select Add Mirror from the Action list.
4. Select a port.
5. Select the name of an ACL from the ACL list.
6. Click Apply.

Figure 200: Configuring ACL Mirroring

The screenshot shows the 'Security > ACL' configuration page. The 'Step' dropdown is set to '3. Configure Interface' and the 'Action' dropdown is set to 'Add Mirror'. Below these, there are two dropdown menus: 'Port' set to '1' and 'ACL' set to 'ips'. At the bottom right, there are 'Apply' and 'Revert' buttons.

To show the ACLs to be mirrored:

1. Select Configure Interface from the Step list.
2. Select Show Mirror from the Action list.
3. Select a port.

Figure 201: Showing the VLANs to Mirror

The screenshot shows the 'Security > ACL' configuration page with the 'Action' dropdown set to 'Show Mirror'. The 'Port' dropdown is set to '1'. Below the dropdowns, there is a table titled 'ACL Mirror List Total: 1'. The table has two columns: a checkbox column and a 'Source Access-List' column. The first row has a checked checkbox and the value 'ips'. At the bottom right, there are 'Delete' and 'Revert' buttons.

	Source Access-List
<input checked="" type="checkbox"/>	ips

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:

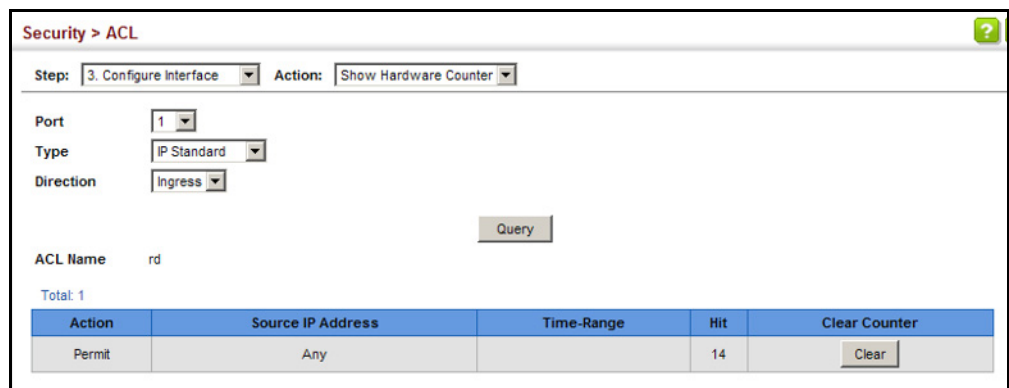
- ◆ **Port** – Port identifier. (Range: 1-12/28)
- ◆ **Type** – Selects the type of ACL.
- ◆ **Direction** – Displays statistics for ingress traffic.
- ◆ **ACL Name** – The ACL bound this port.
- ◆ **Action** – Shows if action is to permit or deny specified packets.
- ◆ *Rules* – Shows the rules for the ACL bound to this port.
- ◆ **Time-Range** – Name of a time range.
- ◆ **Hits** – 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 202: Showing ACL Statistics



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 348](#)). This database is built by DHCP snooping if it is enabled 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 306](#)).

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: 5)
- ◆ **Log Interval** – The interval at which log messages are sent. (Range: 0-86400 seconds; Default: 1 second)

Web Interface

To configure global settings for ARP Inspection:

1. Click Security, ARP Inspection.
2. Select Configure General from the Step list.
3. Enable ARP inspection globally, enable any of the address validation options, and adjust any of the logging parameters if required.
4. Click Apply.

Figure 203: Configuring Global Settings for ARP Inspection

The screenshot shows the configuration page for ARP Inspection. The breadcrumb is "Security > ARP Inspection". The "Step" dropdown is set to "1. Configure General". The configuration options are:

- ARP Inspection Status: Enabled
- ARP Inspection Validation: Dst-MAC IP Src-MAC
- Log Message Number (0-256):
- Log Interval (0-86400): sec

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

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 306](#)).
- ◆ 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:

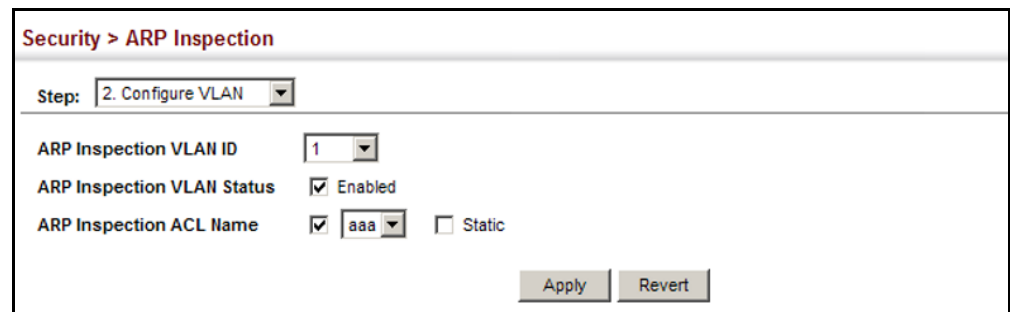
- ◆ **ARP Inspection VLAN ID** – Selects any configured VLAN. (Default: 1)
- ◆ **ARP Inspection VLAN Status** – Enables ARP Inspection for the selected VLAN. (Default: Disabled)
- ◆ **ARP Inspection ACL Name**
 - *ARP ACL* – Allows selection of any configured ARP ACLs. (Default: None)
 - **Static** – When an ARP ACL is selected, and static mode also selected, the switch only performs ARP Inspection and bypasses validation against the DHCP Snooping Bindings database. When an ARP ACL is selected, but static mode is not selected, the switch first performs ARP Inspection and then validation against the DHCP Snooping Bindings database. (Default: Disabled)

Web Interface

To configure VLAN settings for ARP Inspection:

1. Click Security, ARP Inspection.
2. Select Configure VLAN from the Step list.
3. Enable ARP inspection for the required VLANs, select an ARP ACL filter to check for configured addresses, and select the Static option to bypass checking the DHCP snooping bindings database if required.
4. Click Apply.

Figure 204: Configuring VLAN Settings for ARP Inspection



The screenshot shows the 'Security > ARP Inspection' configuration page. At the top, there is a breadcrumb trail and a 'Step:' dropdown menu set to '2. Configure VLAN'. Below this, there are three configuration fields: 'ARP Inspection VLAN ID' with a dropdown menu showing '1', 'ARP Inspection VLAN Status' with a checked checkbox and the text 'Enabled', and 'ARP Inspection ACL Name' with a checked checkbox and a dropdown menu showing 'aaa'. To the right of the 'aaa' dropdown is an unchecked checkbox labeled 'Static'. At the bottom right of the form are two buttons: 'Apply' and '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:

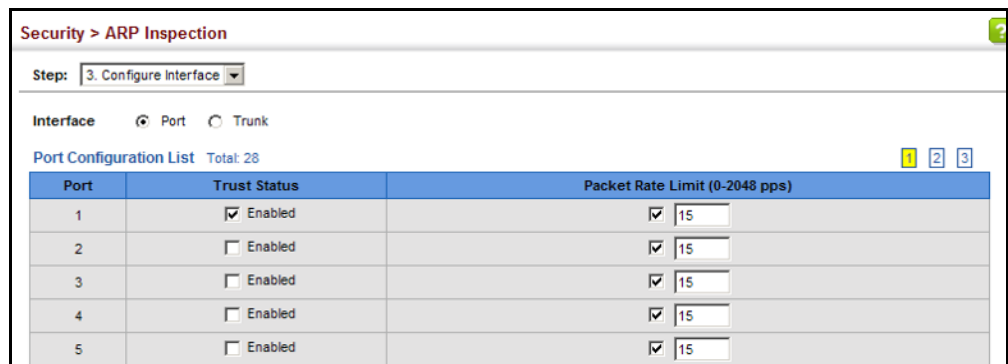
- ◆ **Interface** – Port or trunk identifier.
- ◆ **Trust Status** – Configures the port as trusted or untrusted. (Default: Untrusted)
By default, all untrusted ports are subject to ARP packet rate limiting, and all trusted ports are exempt from ARP packet rate limiting.
Packets arriving on trusted interfaces 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 205: Configuring Interface Settings for ARP Inspection



Displaying ARP Inspection Statistics

Use the Security > ARP Inspection (Show Information - Show Statistics) page to display statistics about the number of ARP packets processed, or dropped for various reasons.

Parameters

These parameters are displayed:

Table 20: ARP Inspection Statistics

Parameter	Description
Received ARP packets before ARP inspection rate limit	Count of ARP packets received but not exceeding the ARP Inspection rate limit.
Dropped ARP packets in the process of ARP inspection rate limit	Count of ARP packets exceeding (and dropped by) ARP rate limiting.
ARP packets dropped by additional validation (IP)	Count of ARP packets that failed the IP address test.
ARP packets dropped by additional validation (Dst-MAC)	Count of packets that failed the destination MAC address test.
Total ARP packets processed by ARP inspection	Count of all ARP packets processed by the ARP Inspection engine.

Table 20: ARP Inspection Statistics (Continued)

Parameter	Description
ARP packets dropped by additional validation (Src-MAC)	Count of packets that failed the source MAC address test.
ARP packets dropped by ARP ACLs	Count of ARP packets that failed validation against ARP ACL rules.
ARP packets dropped by DHCP snooping	Count of packets that failed validation against the DHCP Snooping Binding database.

Web Interface

To display statistics for ARP Inspection:

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

Figure 206: Displaying Statistics for ARP Inspection

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

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

Displaying the ARP Inspection Log

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

Parameters

These parameters are displayed:

Table 21: ARP Inspection Log

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

Table 21: ARP Inspection Log (Continued)

Parameter	Description
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 207: Displaying the ARP Inspection Log

The screenshot shows the 'Security > ARP Inspection' page. At the top, there are two dropdown menus: 'Step: 4. Show Information' and 'Action: Show Log'. Below these is a table titled 'ARP Inspection Log List' with a 'Total: 2' indicator. The table has six columns: VLAN ID, Port, Src. IP Address, Dst. IP Address, Src. MAC Address, and Dst. MAC Address. It contains two rows of data.

VLAN ID	Port	Src. IP Address	Dst. IP Address	Src. MAC Address	Dst. MAC Address
1	15	192.168.1.1	192.168.1.5	11-22-33-44-55-66	AA-BB-CC-DD-EE-FF
1	17	192.168.1.3	192.168.1.23	11-4E-33-75-55-BB	A0-3B-C9-DD-4E-1F

Filtering IP Addresses for Management Access

Use the Security > IP Filter page to create a list of up to 15 IP addresses or IP address groups that are allowed management access to the switch through the web interface, SNMP, or Telnet.

Command Usage

- ◆ The management interfaces are open to all IP addresses by default. Once you add an entry to a filter list, access to that interface is restricted to the specified addresses.
- ◆ If anyone tries to access a management interface on the switch from an invalid address, the switch will reject the connection, enter an event message in the system log, and send a trap message to the trap manager.
- ◆ IP address can be configured for SNMP, web and Telnet access respectively. Each of these groups can include up to five different sets of addresses, either individual addresses or address ranges.

- ◆ When entering addresses for the same group (i.e., SNMP, web or Telnet), the switch will not accept overlapping address ranges. When entering addresses for different groups, the switch will accept overlapping address ranges.
- ◆ You cannot delete an individual address from a specified range. You must delete the entire range, and reenter the addresses.
- ◆ You can delete an address range just by specifying the start address, or by specifying both the start address and end address.

Parameters

These parameters are displayed:

- ◆ **Mode**
 - **Web** – Configures IP address(es) for the web group.
 - **SNMP** – Configures IP address(es) for the SNMP group.
 - **Telnet** – Configures IP address(es) for the Telnet group.
- ◆ **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).
4. Enter the IP addresses or range of addresses that are allowed management access to an interface.
5. Click Apply

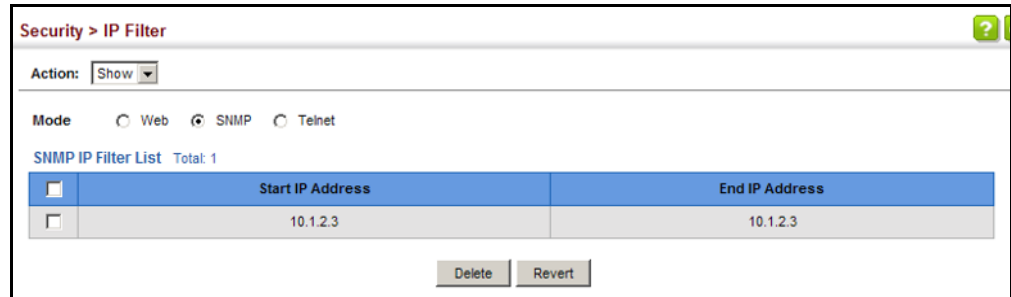
Figure 208: Creating an IP Address Filter for Management Access

The screenshot shows a web interface for configuring an IP filter. At the top, it says "Security > IP Filter". Below that, there is an "Action:" label followed by a dropdown menu currently set to "Add". Underneath, there is a "Mode" section with three radio buttons: "Web" (which is selected), "SNMP", and "Telnet". Below the mode selection, there are two input fields: "Start IP Address" which contains the text "10.1.2.3", and "End IP Address" which is currently empty. At the bottom right of the form, there are two buttons: "Apply" and "Revert".

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

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

Figure 209: Showing IP Addresses Authorized for Management Access



Configuring Port Security

Use the Security > Port Security page to configure the maximum number of device MAC addresses that can be learned by a switch port, stored in the address table, and authorized to access the network.

When port security is enabled on a port, the switch stops learning new MAC addresses on the specified port when it has reached a configured maximum number. Only incoming traffic with source addresses already stored in the address table will be authorized to access the network through that port. If a device with an unauthorized MAC address attempts to use the switch port, the intrusion will be detected and the switch can automatically take action by disabling the port and sending a trap message.

Command Usage

- ◆ The default maximum number of MAC addresses allowed on a secure port is zero (that is, disabled). To use port security, you must configure the maximum number of addresses allowed on a port.
- ◆ To configure the maximum number of address entries which can be learned on a port, 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 165](#)).

- ◆ 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 95](#)).
- ◆ 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.

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 - 256, 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 276](#).

- ◆ **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 210: Configuring Port Security

Port	Security Status	Port Status	Action	Max MAC Count (0-256)	Current MAC Count	MAC Filter	MAC Filter ID	Last Intrusion MAC	Last Time Detected Intrusion MAC
1	<input type="checkbox"/> Enabled	Secure/Down	None	0	3	Disabled	0	NA	NA
2	<input type="checkbox"/> Enabled	Secure/Down	None	0	0	Disabled	0	NA	NA
3	<input type="checkbox"/> Enabled	Secure/Down	None	0	0	Disabled	0	NA	NA

Configuring 802.1X Port Authentication

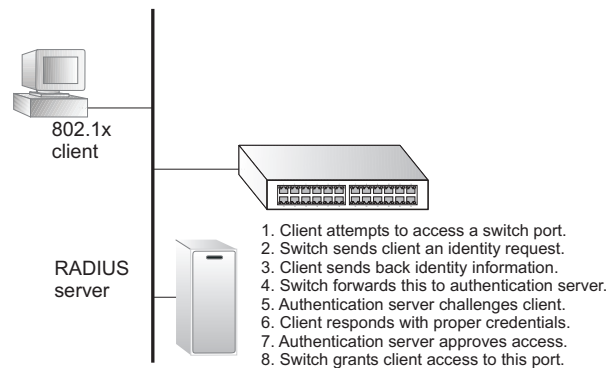
Network switches can provide open and easy access to network resources by simply attaching a client PC. Although this automatic configuration and access is a desirable feature, it also allows unauthorized personnel to easily intrude and possibly gain access to sensitive network data.

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 211: Configuring Port Security



The operation of 802.1X on the switch requires the following:

- ◆ The switch must have an IP address assigned.
- ◆ RADIUS authentication must be enabled on the switch and the IP address of the RADIUS server specified.
- ◆ 802.1X must be enabled globally for the switch.
- ◆ Each switch port that will be used must be set to dot1X “Auto” mode.
- ◆ Each client that needs to be authenticated must have dot1X client software installed and properly configured.
- ◆ The RADIUS server and 802.1X client support EAP. (The switch only supports EAPOL in order to pass the EAP packets from the server to the client.)
- ◆ The RADIUS server and client also have to support the same EAP authentication type – MD5, PEAP, TLS, or TTLS. (Native support for these encryption methods is

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

Configuring 802.1X Global Settings

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

Parameters

These parameters are displayed:

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

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

When this device is functioning as an edge switch but does not require any attached clients to be authenticated, **EAPOL Pass Through** can be disabled to discard unnecessary EAPOL traffic.
- ◆ **Identity Profile User Name** – The dot1x supplicant user name. (Range: 1-8 characters)

The global supplicant user name and password are used to identify this switch as a supplicant when responding to an MD5 challenge from the authenticator. These parameters must be set when this switch passes client authentication requests to another authenticator on the network (see [“Configuring Port Supplicant Settings for 802.1X”](#) on page 330).
- ◆ **Set Password** – Allows the dot1x supplicant password to be entered.
- ◆ **Identity Profile Password** – The dot1x supplicant password used to identify this switch as a supplicant when responding to an MD5 challenge from the authenticator. (Range: 1-8 characters)
- ◆ **Confirm Profile Password** – This field is used to confirm the dot1x supplicant password.
- ◆ **Default** – Sets all configurable 802.1X global and port settings to their default values.

Web Interface

To configure global settings for 802.1X:

1. Click Security, Port Authentication.
2. Select Configure Global from the Step list.
3. Enable 802.1X globally for the switch, and configure EAPOL Pass Through if required. Then set the user name and password to use when the switch responds an MD5 challenge from the authentication server.
4. Click Apply

Figure 212: Configuring Global Settings for 802.1X Port Authentication

Security > Port Authentication

Step: 1. Configure Global

System Authentication Control Enabled

EAPOL Pass Through Enabled

Identity Profile User Name

Set Password

Identity Profile Password

Confirm Profile Password

Apply Revert

Default [Click this button to set 802.1X global/port settings to default values.](#)

Configuring Port Authenticator Settings for 802.1X

Use the Security > Port Authentication (Configure Interface – Authenticator) page to configure 802.1X port settings for the switch as the local authenticator. When 802.1X is enabled, you need to configure the parameters for the authentication process that runs between the client and the switch (i.e., authenticator), as well as the client identity lookup process that runs between the switch and authentication server.

Command Usage

- ◆ When the switch functions as a local authenticator between supplicant devices attached to the switch and the authentication server, configure the parameters for the exchange of EAP messages between the authenticator and clients on the Authenticator configuration page.
- ◆ When devices attached to a port must submit requests to another authenticator on the network, configure the Identity Profile parameters on the Configure Global page (see [“Configuring 802.1X Global Settings” on page 325](#)) which identify this switch as a supplicant, and configure the supplicant parameters for those ports which must authenticate clients through the

remote authenticator (see [“Configuring Port Supplicant Settings for 802.1X” on page 330](#)).

- ◆ 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.
- ◆ **Supplicant** – Indicates the MAC address of a connected client.
- ◆ **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: 1-256; 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 252.](#))
- ◆ **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 138](#)) and mapped on each port (See [“Configuring Network Access for Ports” on page 273](#)).

Supplicant List

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

Authenticator PAE State Machine

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

Backend State Machine

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

Reauthentication State Machine

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

Web Interface

To configure port authenticator settings for 802.1X:

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

5. Click Apply

Figure 213: Configuring Interface Settings for 802.1X Port Authenticator

Security > Port Authentication

Step: 2. Configure Interface

Type Authenticator Supplicant

Port: 4

Status: Enabled

Authorized: Yes

Control Mode: Auto

Operation Mode: Single-Host

Max Count (1-256): 5

Max Request (1-10): 2

Quiet Period (1-65535): 60 sec

Tx Period (1-65535): 30 sec

Supplicant Timeout (1-65535): 30 sec

Server Timeout: 10 sec

Re-authentication Status: Enabled

Re-authentication Period (1-65535): 3600 sec

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

Intrusion Action: Block Traffic

Supplicant List Total: 1

Supplicant	Authenticator PAE State Machine			Backend State Machine			Reauthentication State Machine
	State	Reauth Count	Current Identifier	State	Request Count	Identifier (Server)	State
00-E0-29-94-34-65	Authenticated	0	10	Idle	0	9	Initialize

Apply Revert

Configuring Port Supplicant Settings for 802.1X

Use the Security > Port Authentication (Configure Interface – Supplicant) page to configure 802.1X port settings for supplicant requests issued from a port to an authenticator on another device. When 802.1X is enabled and the control mode is set to Force-Authorized (see “Configuring Port Authenticator Settings for 802.1X” on page 326), you need to configure the parameters for the client supplicant process if the client must be authenticated through another device in the network.

Command Usage

- ◆ When devices attached to a port must submit requests to another authenticator on the network, configure the Identity Profile parameters on the Configure Global page (see “Configuring 802.1X Global Settings” on page 325) which identify this switch as a supplicant, and configure the supplicant parameters for those ports which must authenticate clients through the remote authenticator on this configuration page. When PAE supplicant mode is enabled on a port, it will not respond to dot1x messages meant for an authenticator.

- ◆ This switch can be configured to serve as the authenticator on selected ports by setting the Control Mode to Auto on the Authenticator configuration page, and as a supplicant on other ports by the setting the control mode to Force-Authorized on that configuration page and enabling the PAE supplicant on the Supplicant configuration page.

Parameters

These parameters are displayed:

- ◆ **Port** – Port number.
- ◆ **PAE Supplicant** – Enables PAE supplicant mode. (Default: Disabled)
If the attached client must be authenticated through another device in the network, supplicant status must be enabled.
Supplicant status can only be enabled if PAE Control Mode is set to “Force-Authorized” on this port (see [“Configuring Port Authenticator Settings for 802.1X” on page 326](#)).
PAE supplicant status cannot be enabled if a port is a member of trunk or LACP is enabled on the port.
- ◆ **Authentication Period** – The time that a supplicant port waits after the Maximum Start count (described below) has been exceeded for a response from the authenticator. (Range: 1-65535 seconds; Default: 30 seconds)
- ◆ **Held Period** – The time that a supplicant port waits before resending its credentials to find a new an authenticator. (Range: 1-65535 seconds; Default: 60 seconds)
- ◆ **Start Period** – The time that a supplicant port waits before resending an EAPOL start frame to the authenticator. (Range: 1-65535 seconds; Default: 30 seconds)
- ◆ **Maximum Start** – The maximum number of times that a port supplicant will send an EAP start frame to the client before assuming that the client is 802.1X unaware. (Range: 1-65535; Default: 3)
- ◆ **Authenticated** – Shows whether or not the supplicant has been authenticated.

Web Interface

To configure port authenticator settings for 802.1X:

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

Figure 214: Configuring Interface Settings for 802.1X Port Supplicant

The screenshot shows the 'Security > Port Authentication' configuration page. The 'Step' dropdown is set to '2. Configure Interface'. Under the 'Type' section, 'Supplicant' is selected with a radio button. The 'Port' is set to '2'. The 'PAE Supplicant' checkbox is checked and labeled 'Enabled'. The 'Authentication Period (1-65535)' is set to '30', 'Held Period (1-65535)' is '60', 'Start Period (1-65535)' is '30', and 'Maximum Start (1-65535)' is '3'. The 'Authenticated' status is 'No'. At the bottom right, there are 'Apply' and 'Revert' buttons.

Displaying 802.1X Statistics Use the Security > Port Authentication (Show Statistics) page to display statistics for dot1x protocol exchanges for any port.

Parameters

These parameters are displayed:

Table 22: 802.1X Statistics

Parameter	Description
<i>Authenticator</i>	
Rx EAPOL Start	The number of EAPOL Start frames that have been received by this Authenticator.
Rx EAPOL Logoff	The number of EAPOL Logoff frames that have been received by this Authenticator.
Rx EAPOL Invalid	The number of EAPOL frames that have been received by this Authenticator in which the frame type is not recognized.

Table 22: 802.1X Statistics (Continued)

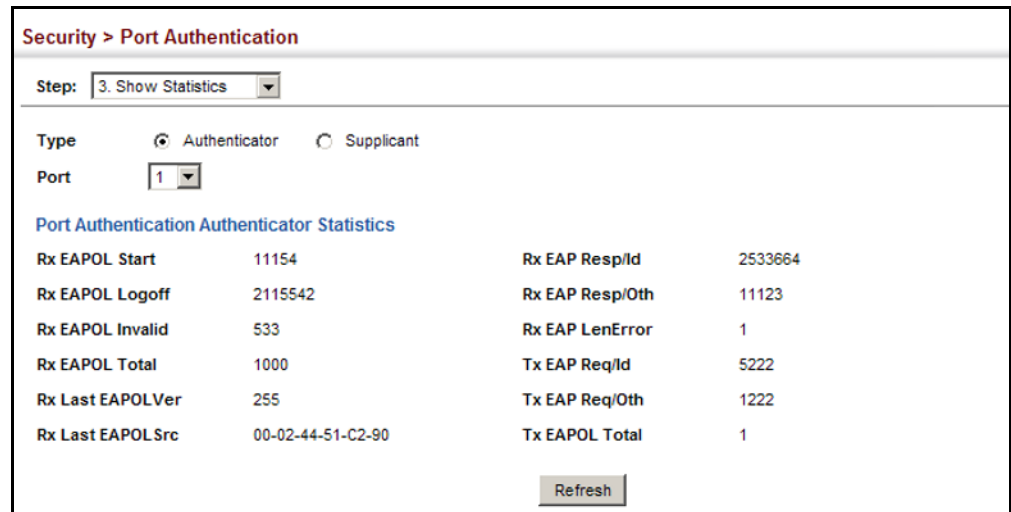
Parameter	Description
Rx EAPOL Total	The number of valid EAPOL frames of any type that have been received by this Authenticator.
Rx Last EAPOLVer	The protocol version number carried in the most recent EAPOL frame received by this Authenticator.
Rx Last EAPOLSrc	The source MAC address carried in the most recent EAPOL frame received by this Authenticator.
Rx EAP Resp/Id	The number of EAP Resp/Id frames that have been received by this Authenticator.
Rx EAP Resp/Oth	The number of valid EAP Response frames (other than Resp/Id frames) that have been received by this Authenticator.
Rx EAP LenError	The number of EAPOL frames that have been received by this Authenticator in which the Packet Body Length field is invalid.
Tx EAP Req/Id	The number of EAP Req/Id frames that have been transmitted by this Authenticator.
Tx EAP Req/Oth	The number of EAP Request frames (other than Rq/Id frames) that have been transmitted by this Authenticator.
Tx EAPOL Total	The number of EAPOL frames of any type that have been transmitted by this Authenticator.
<i>Supplicant</i>	
Rx EAPOL Invalid	The number of EAPOL frames that have been received by this Supplicant in which the frame type is not recognized.
Rx EAPOL Total	The number of valid EAPOL frames of any type that have been received by this Supplicant.
Rx Last EAPOLVer	The protocol version number carried in the most recent EAPOL frame received by this Supplicant.
Rx Last EAPOLSrc	The source MAC address carried in the most recent EAPOL frame received by this Supplicant.
Rx EAP Resp/Id	The number of EAP Resp/Id frames that have been received by this Supplicant.
Rx EAP Resp/Oth	The number of valid EAP Response frames (other than Resp/Id frames) that have been received by this Supplicant.
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 Rq/Id frames) that have been transmitted by this Supplicant.

Web Interface

To display port authenticator statistics for 802.1X:

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

Figure 215: Showing Statistics for 802.1X Port Authenticator



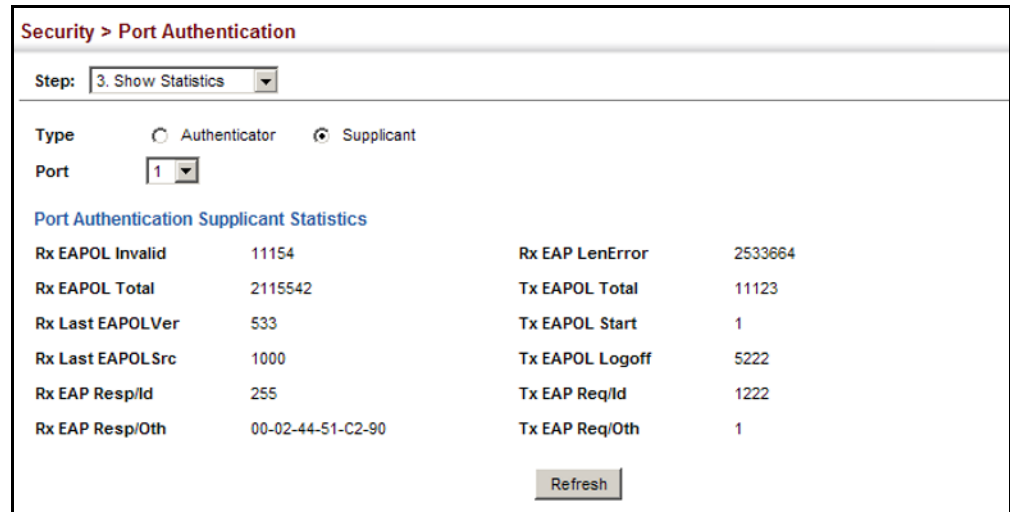
The screenshot shows the 'Security > Port Authentication' web interface. At the top, there is a breadcrumb trail 'Security > Port Authentication'. Below it, a 'Step:' dropdown menu is set to '3. Show Statistics'. Under the 'Type' section, the 'Authenticator' radio button is selected, and the 'Supplicant' radio button is unselected. The 'Port' dropdown menu is set to '1'. The main content area is titled 'Port Authentication Authenticator Statistics' and displays a table of statistics. A 'Refresh' button is located at the bottom right of the statistics table.

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

To display port supplicant statistics for 802.1X:

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

Figure 216: Showing Statistics for 802.1X Port Supplicant



DoS Protection

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

Command Usage

- ◆ All DoS functions are disabled by default.

Parameters

These parameters are displayed:

Global Protection

- ◆ **DoS Protection** – Enables DoS protection globally on the switch.

Protection for IPv4

- ◆ **Invalid IP Address** – Protects against attacks in which the source IP address and the destination IP address are the same.
- ◆ **Invalid Header Length** – Protects against attacks which send IP packets with an incorrect header length or IP data length. Such packets are not allowed by the system, but their abundant number can cause computer crashes and other system errors.

These packets may have any of the following attributes:

- ◆ Header length is less than 4 bytes
- ◆ Raw IP data length is less than header length * 4
- ◆ **Invalid Source IP Address** – Protects against attacks in which hackers replace the source address in packets sent to the victim with an invalid source IP address to protect the identity of the sender or to mislead the receiver as to the origin and validity of sent data. These attacks may send a constant stream of packets with an invalid source address such as 127.0.0.1, causing receiver to respond in the desired manner, while continuing to hide the identity of the attacker. This type of attack is especially effective since the packets may seem to come from different sources and thus making the perpetrators hard to trace.

These packets may have any of the following attributes:

- ◆ $224.0.0.0 \leq \text{source IP address} \leq 240.0.0.0$
- ◆ Source IP address is 127.*.*
- ◆ Source IP address is 255.255.255.255

One of the main reasons for forging a source address while staging a DoS attack is to avoid detection upon staging the attack. The other reason is to stage a twofold attack. One example of such an attack is a smurf attack. In a smurf attack, the attacker attacks in two places at the same time. Not only is the end target affected by the large number of echo replies received, but the network that acts as the reflector is also affected by the large amount of traffic.

- ◆ **Invalid Destination IP Address** – Protects against invalid IP destination address attacks. When a stream such packets are received, this can indicate a denial-of-service (DoS) attempt or just a packet generator using RAW sockets on the network.

These packets may have any of the following attributes:

- ◆ Destination IP address is 127.*.*
- ◆ Destination IP address is 0.0.0.0

Protection for IPv6

- ◆ **Invalid IP Address** – Protects against attacks in which the source IP address and the destination IP address are the same.
- ◆ **Invalid Header Length** – Protects against attacks which send IP packets with an incorrect header length. Such packets are not allowed by the system, but their abundant number can cause computer crashes and other system errors.
- ◆ **Invalid Source IP Address** – Protects against attacks in which hackers replace the source address in packets sent to the victim with an invalid source IPv6 address to protect the identity of the sender or to mislead the receiver as to the origin and validity of sent data. These attacks may send a constant stream of packets with an invalid source address, causing receiver to respond in the desired manner, while continuing to hide the identity of the attacker. This type

of attack is especially effective since the packets seem to come from different sources and thus making the perpetrators hard to trace.

These packets may have any of the following attributes:

- ◆ Source IP address is ::1
- ◆ Source IP address is 0xFF00::/8
- ◆ **Invalid Destination IP Address** – Protects against invalid IPv6 destination address attacks. When a stream such packets are received, this can indicate a denial-of-service (DoS) attempt or just a packet generator using RAW sockets on the network.

Protection for TCP

- ◆ **Invalid Header Length** – Protects against attacks which send TCP packets with an incorrect header length. Such packets are not allowed by the system, but their abundant number can cause computer crashes and other system errors. In these packets, the TCP raw header size is less than the minimum size defined for a TCP header (i.e., the data offset < 5).
- ◆ **Blat Block** – Protects against attacks in which a specially crafted packet is sent to a host where the source host port is the same as the destination host port. The system attempts to reply to itself, resulting in system lockup.
- ◆ **SYN/URG Block** – Protects against attacks in which a TCP SYN/URG message is used to interrupt or abort the queued stream. In these packets, SYN=1 and URG=1.
- ◆ **SYN/PSH Block** – Protects against attacks in which a TCP SYN/PSH message is used to force the TCP stack to send this data immediately up to the receiving application. In these packets, SYN=1 and PSH=1
- ◆ **SYN/ACK/PSH Block** – Protects against attacks in which a TCP SYN/ACK/PSH message sequence is used to cause problems for some operating systems which do not acknowledge this as a valid sequence. In these packets, SYN=1, ACK=1 and PSH=1.
- ◆ **Xmas Scan** – Protects against Xmas-scan in which 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. In these packets, FIN=1, URG= 1 and PSH = 1.
- ◆ **NULL Scan** – Protects against null-scan attacks in which 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.

- ◆ **SYN/FIN Scan** – Protects against SYN/FIN-scan attacks in which 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.

In these packets, SYN=1 and FIN=1.

- ◆ **SYN/RST Scan** – Protects against SYN/RST-scan attacks in which a TCP SYN/RST scan message is used to stop an ongoing TCP session. An attacker can forge a set of Synchronize (SYN) and Reset (RST) packets in an attempt to guess a TCP sequence number within a narrow range (or TCP window) of values. Successful exploitation of this issue results in a termination of the TCP session. Depending on the targeted software or hardware, the outcome may result in a simple denial of service, or it may leave the system in an unpredictable state, possibly leading to data loss or additional vulnerabilities.

In these packets, SYN=1 and RST=1.

- ◆ **SYN Flood** – Protects against flooding attacks in which a perpetrator sends a succession of TCP synchronization requests (with or without a spoofed source IP address) to a target and never returns ACK packets. These half-open connections will bind up resources on the target, and no new connections can be made, resulting in denial of service. (Maximum allowed rate: 64-2048 kbits/second)

In these packets, SYN=1.

Protection for UDP

- ◆ **Invalid Header Length** – Protects against attacks which send UDP packets with an incorrect header length. Such packets are not allowed by the system, but their abundant number can cause computer crashes and other system errors.

In these packets, the UDP raw data length is less than 8 bytes.

- ◆ **Blat Block** – Protects against attacks in which a specially crafted packet is sent to a host where the source host port is the same as the destination host port. The system attempts to reply to itself, resulting in system lockup.

- ◆ **Flood** – Protects against flooding attacks in which a perpetrator sends a large number of UDP packets (with or without a spoofed source IP address) to random ports on a remote host. The target will determine that an application is listening at that port, and reply with an ICMP "Destination Unreachable" packet. It will be forced to send many ICMP packets, eventually leading it to be unreachable by other clients. (Maximum allowed rate: 64-2048 kbits/second)

Protection for ICMP

- ◆ **Smurf** – Protects against smurf attacks in which a perpetrator generates a large amount of spoofed ICMP Echo Request traffic to the broadcast destination IP address (255.255.255.255), all of which uses a spoofed source address of the intended victim. The victim should crash due to the many interrupts required to send ICMP Echo response packets.
- ◆ **Ping of Death** – Protects against ping-of-death attacks in which an attacker deliberately sends an IP packet larger than the 65,536 bytes allowed by the IP protocol, or by using fragmentation in which a packet broken down into fragments could add up to more than the allowed 65,536 bytes. Some operating systems do not know how to respond when they receive an oversized packet, so they freeze, crash, or reboot.
- ◆ **Nuke** – Protects against nuke attacks which send fragmented or otherwise invalid ICMP packets using a modified ping utility to repeatedly send the corrupted data, thus slowing down the affected host until it comes to a complete stop. Nuke attacks may also send an ICMP packets (usually through port 139) with a “destination unreachable” message to cause connection breaks.
- ◆ **Flood** – Protects against flooding attacks in which large amounts of (or just over-sized) ICMP packets are sent to a host in order to attempt to crash the TCP/IP stack on the host. An ICMP flood can consist of any type of ICMP message, including smurf, ping-flood, or ping-of-death attacks. (Maximum allowed rate: 64-2048 kbits/second)

Other Protection Commands

- ◆ **Echo/Chargen** – Protects against Echo/Chargen attacks in which the echo service repeats anything sent to it, and the chargen (character generator) service generates a continuous stream of data. When used together, they create an infinite loop and result in denial-of-service. (Maximum allowed rate: 64-2048 kbits/second)

Web Interface

To protect against DoS attacks:

1. Click Security, DoS Protection.
2. Enable any of the required protection features, and set the maximum rate of ingress for flooding attacks or Echo/Chargen.
3. Click Apply

Figure 217: Configuring DoS Protection

Security > DoS Protection

DoS Protection

DoS Protection Enabled

DoS Protection IP

Invalid IP Address Enabled

Invalid Header Length Enabled

Invalid Source IP Address Enabled

Invalid Destination IP Address Enabled

DoS Protection IPv6

Invalid IP Address Enabled

Invalid Header Length Enabled

Invalid Source IP Address Enabled

Invalid Destination IP Address Enabled

DoS Protection TCP

Invalid Header Length Enabled

Blat Block Enabled

SYN-URG Block Enabled

SYN-PSH Block Enabled

SYN-ACK-PSH Block Enabled

XMAS Scan Enabled

NULL Scan Enabled

SYN-FIN Scan Enabled

SYN-RST Scan Enabled

SYN Flood (64-2048) kbps

DoS Protection UDP

Invalid Header Length Enabled

Blat Block Enabled

Flood (64-2048) kbps

DoS Protection ICMP

Smurf Enabled

Ping of Death Enabled

Nuke Enabled

Flood (64-2048) kbps

DoS Protection Echo/Chargen

DoS Protection Echo/Chargen (64-2048) kbps

Apply Revert

IP Source Guard

IP Source Guard is a security feature that filters IP traffic on network interfaces based on manually configured entries in the IP Source Guard table, or dynamic entries in the DHCP Snooping table when enabled (see [“DHCP Snooping” on page 346](#)). IP source guard can be used to prevent traffic attacks caused when a host tries to use the IP address of a neighbor to access the network. This section describes commands used to configure IP Source Guard.

Configuring Ports for IP Source Guard Use the Security > IP Source Guard > Port Configuration page to set the filtering type based on source IP address, or source IP address and MAC address pairs.

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

Command Usage

- ◆ Setting source guard mode to SIP (Source IP) or SIP-MAC (Source IP and MAC) enables this function on the selected port. Use the SIP option to check the VLAN ID, source IP address, and port number against all entries in the binding table. Use the SIP-MAC option to check these same parameters, plus the source MAC address. If no matching entry is found, the packet is dropped.



Note: Multicast addresses cannot be used by IP Source Guard.

- ◆ When enabled, traffic is filtered based upon dynamic entries learned via DHCP snooping (see [“DHCP Snooping” on page 346](#)), 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.
- ◆ Filtering rules are implemented as follows:
 - If DHCP snooping is disabled (see [page 348](#)), IP source guard will check the VLAN ID, source IP address, port number, and source MAC address (for the SIP-MAC option). If a matching entry is found in the binding table and the entry type is static IP source guard binding, the packet will be forwarded.
 - If DHCP snooping is enabled, IP source guard will check the VLAN ID, source IP address, port number, and source MAC address (for the SIP-MAC option). If a matching entry is found in the binding table and the entry type is static IP source guard binding, or dynamic DHCP snooping binding, the packet will be forwarded.

- If IP source guard is enabled on an interface for which IP source bindings have not yet been configured (neither by static configuration in the IP source guard binding table nor dynamically learned from DHCP snooping), the switch will drop all IP traffic on that port, except for DHCP packets.

Parameters

These parameters are displayed:

- ◆ **Filter Type** – Configures the switch to filter inbound traffic based on source IP address, or source IP address and corresponding MAC address. (Default: None)
 - **None** – Disables IP source guard filtering on the port.
 - **SIP** – Enables traffic filtering based on IP addresses stored in the binding table.
 - **SIP-MAC** – Enables traffic filtering based on IP addresses and corresponding MAC addresses stored in the binding table.
- ◆ **Max Binding Entry** – The maximum number of entries that can be bound to an interface. (Range: 1-5; Default: 5)

This parameter sets the maximum number of address entries that can be mapped to an interface in the binding table, including both dynamic entries discovered by DHCP snooping (see “[DHCP Snooping](#)” on page 346) and static entries set by IP source guard (see “[Configuring Static Bindings for IP Source Guard](#)” on page 343).

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 218: Setting the Filter Type for IP Source Guard

Port	Filter Type	Max Binding Entry (1-5)
1	None	5
2	None	5
3	None	5
4	None	5
5	SIP	3

Configuring Static Bindings for IP Source Guard

Use the Security > IP Source Guard > Static Configuration page 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

- ◆ Static addresses entered in the source guard binding table are automatically configured with an infinite lease time. Dynamic entries learned via DHCP snooping are configured by the DHCP server itself.
- ◆ Static bindings are processed as follows:
 - 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.
 - 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.
- ◆ **IP Address** – A valid unicast IP address, including classful types A, B or C.

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.
- ◆ **IP Address** – IP address corresponding to the client.
- ◆ **Lease Time** – The time for which this IP address is leased to the client. (This value is zero for all static addresses.)

Web Interface

To configure static bindings for IP Source Guard:

1. Click Security, IP Source Guard, Static Configuration.
2. Select Add from the Action list.
3. Enter the required bindings for each port.
4. Click Apply

Figure 219: Configuring Static Bindings for IP Source Guard

Security > IP Source Guard > Static Binding

Action: Add

Port: 1

VLAN: 1

MAC Address: 00-10-B5-F4-00-01

IP Address: 10.2.44.96

Apply Revert

To display static bindings for IP Source Guard:

1. Click Security, IP Source Guard, Static Configuration.
2. Select Show from the Action list.

Figure 220: Displaying Static Bindings for IP Source Guard

Security > IP Source Guard > Static Binding

Action: Show

Static Binding List Total: 1

<input type="checkbox"/>	VLAN	MAC Address	Interface	IP Address	Lease Time (sec)
<input type="checkbox"/>	1	00-10-B5-F4-00-01	Unit 1 / Port 1	192.168.0.23	0

Delete Revert

Displaying Information for Dynamic IP Source Guard Bindings

Use the Security > IP Source Guard > Dynamic Binding page to display the source-guard binding table for a selected interface.

Parameters

These parameters are displayed:

Query by

- ◆ **Port** – A port on this switch.

- ◆ **VLAN** – ID of a configured VLAN (Range: 1-4094)
- ◆ **MAC Address** – A valid unicast MAC address.
- ◆ **IP Address** – A valid unicast IP address, including classful types A, B or C.

Dynamic Binding List

- ◆ **VLAN** – VLAN to which this entry is bound.
- ◆ **MAC Address** – Physical address associated with the entry.
- ◆ **Interface** – Port to which this entry is bound.
- ◆ **IP Address** – IP address corresponding to the client.
- ◆ **Lease Time** – The time for which this IP address is leased to the client.

Web Interface

To display the binding table for IP Source Guard:

1. Click Security, IP Source Guard, Dynamic Binding.
2. Mark the search criteria, and enter the required values.
3. Click Query

Figure 221: Showing the IP Source Guard Binding Table

VLAN	MAC Address	Interface	IP Address	Lease Time (sec)
1	00-10-B5-F4-00-01	Unit 1 / Port 2	10.2.44.96	5
1	00-10-B5-F4-00-02	Unit 1 / Port 4	10.2.44.97	25
2	00-10-B5-F4-00-03	Unit 1 / Port 7	10.2.44.98	47

DHCP Snooping

The addresses assigned to DHCP clients on insecure ports can be carefully controlled using the dynamic bindings registered with DHCP Snooping (or using the static bindings configured with IP Source Guard). DHCP snooping allows a switch to protect a network from rogue DHCP servers or other devices which send port-related information to a DHCP server. This information can be useful in tracking an IP address back to a physical port.

Command Usage

DHCP Snooping Process

- ◆ Network traffic may be disrupted when malicious DHCP messages are received from an outside source. DHCP snooping is used to filter DHCP messages received on a non-secure interface from outside the network or fire wall. When DHCP snooping is enabled globally and enabled on a VLAN interface, DHCP messages received on an untrusted interface from a device not listed in the DHCP snooping table will be dropped.
- ◆ Table entries are only learned for trusted interfaces. An entry is added or removed dynamically to the DHCP snooping table when a client receives or releases an IP address from a DHCP server. Each entry includes a MAC address, IP address, lease time, VLAN identifier, and port identifier.
- ◆ The rate limit for the number of DHCP messages that can be processed by the switch is 100 packets per second. Any DHCP packets in excess of this limit are dropped.
- ◆ When DHCP snooping is enabled, DHCP messages entering an untrusted interface are filtered based upon dynamic entries learned via DHCP snooping.
- ◆ Filtering rules are implemented as follows:
 - If the global DHCP snooping is disabled, all DHCP packets are forwarded.
 - If DHCP snooping is enabled globally, and also enabled on the VLAN where the DHCP packet is received, all DHCP packets are forwarded for a *trusted* port. If the received packet is a DHCP ACK message, a dynamic DHCP snooping entry is also added to the binding table.
 - If DHCP snooping is enabled globally, and also enabled on the VLAN where the DHCP packet is received, but the port is *not trusted*, it is processed as follows:
 - If the DHCP packet is a reply packet from a DHCP server (including OFFER, ACK or NAK messages), the packet is dropped.
 - If the DHCP packet is from a client, such as a DECLINE or RELEASE message, the switch forwards the packet only if the corresponding entry is found in the binding table.

- If the DHCP packet is from a client, such as a DISCOVER, REQUEST, INFORM, DECLINE or RELEASE message, the packet is forwarded if MAC address verification is disabled. However, if MAC address verification is enabled, then the packet will only be forwarded if the client's hardware address stored in the DHCP packet is the same as the source MAC address in the Ethernet header.
- If the DHCP packet is not a recognizable type, it is dropped.
- If a DHCP packet from a client passes the filtering criteria above, it will only be forwarded to trusted ports in the same VLAN.
- If a DHCP packet is from server is received on a trusted port, it will be forwarded to both trusted and untrusted ports in the same VLAN.
- If the DHCP snooping is globally disabled, all dynamic bindings are removed from the binding table.
- *Additional considerations when the switch itself is a DHCP client* – The port(s) through which the switch submits a client request to the DHCP server must be configured as trusted. Note that the switch will not add a dynamic entry for itself to the binding table when it receives an ACK message from a DHCP server. Also, when the switch sends out DHCP client packets for itself, no filtering takes place. However, when the switch receives any messages from a DHCP server, any packets received from untrusted ports are dropped.

DHCP Snooping Option 82

- ◆ DHCP provides a relay mechanism for sending information about its DHCP clients or the relay agent itself to the DHCP server. Also known as DHCP Option 82, it allows compatible DHCP servers to use the information when assigning IP addresses, or to set other services or policies for clients. It is also an effective tool in preventing malicious network attacks from attached clients on DHCP services, such as IP Spoofing, Client Identifier Spoofing, MAC Address Spoofing, and Address Exhaustion.
- ◆ DHCP Snooping must be enabled for Option 82 information to be inserted into request packets.
- ◆ When the DHCP Snooping Information Option 82 is enabled, the requesting client (or an intermediate relay agent that has used the information fields to describe itself) can be identified in the DHCP request packets forwarded by the switch and in reply packets sent back from the DHCP server. This information may specify the MAC address or IP address of the requesting device (that is, the switch in this context).

By default, the switch also fills in the Option 82 circuit-id field with information indicating the local interface over which the switch received the DHCP client request, including the port and VLAN ID. This allows DHCP client-server exchange messages to be forwarded between the server and client without having to flood them to the entire VLAN.

- ◆ If DHCP Snooping Information Option 82 is enabled on the switch, information may be inserted into a DHCP request packet received over any VLAN (depending on DHCP snooping filtering rules). The information inserted into the relayed packets includes the circuit-id and remote-id, as well as the gateway Internet address.
- ◆ When the switch receives DHCP packets from clients that already include DHCP Option 82 information, the switch can be configured to set the action policy for these packets. The switch can either drop the DHCP packets, keep the existing information, or replace it with the switch's relay information.
- ◆ DHCP Snooping Information Option 82 cannot both be enabled at the same time.

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:

- ◆ **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 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.
- ◆ **DHCP Snooping Information Option Remote ID** – Specifies the MAC address, IP address, or arbitrary identifier of the requesting device (i.e., the switch in this context).
 - **MAC Address** – Inserts a MAC address in the remote ID sub-option for the DHCP snooping agent (i.e., the MAC address of the switch's CPU). This attribute can be encoded in Hexadecimal or ASCII.
 - **IP Address** – Inserts an IP address in the remote ID sub-option for the DHCP snooping agent (i.e., the IP address of the management interface). This attribute can be encoded in Hexadecimal or ASCII.
 - *string* - An arbitrary string inserted into the remote identifier field. (Range: 1-32 characters)

- ◆ **DHCP Snooping Information Option Policy** – Specifies how to handle DHCP client request packets which already contain Option 82 information.
 - **Drop** – Drops the client’s request packet instead of relaying it.
 - **Keep** – Retains the Option 82 information in the client request, and forwards the packets to trusted ports.
 - **Replace** – Replaces the Option 82 information circuit-id and remote-id fields in the client’s request with information about the relay agent itself, inserts the relay agent’s address (when DHCP snooping is enabled), and forwards the packets to trusted ports. (This is the default policy.)

Web Interface

To configure global settings for DHCP Snooping:

1. Click IP Service, DHCP, Snooping.
2. Select Configure Global from the Step list.
3. Select the required options for the general DHCP snooping process and for the DHCP Option 82 information option.
4. Click Apply

Figure 222: Configuring Global Settings for DHCP Snooping

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

Command Usage

- ◆ When DHCP snooping is enabled globally on the switch, and enabled on the specified VLAN, DHCP packet filtering will be performed on any untrusted ports within the VLAN.

- ◆ When the DHCP snooping is globally disabled, DHCP snooping can still be configured for specific VLANs, but the changes will not take effect until DHCP snooping is globally re-enabled.
- ◆ When DHCP snooping is globally enabled, and DHCP snooping is then disabled on a VLAN, all dynamic bindings learned for this VLAN are removed from the binding table.

Parameters

These parameters are displayed:

- ◆ **VLAN** – ID of a configured VLAN. (Range: 1-4094)
- ◆ **DHCP Snooping Status** – Enables or disables DHCP snooping for the selected VLAN. When DHCP snooping is enabled globally on the switch, and enabled on the specified VLAN, DHCP packet filtering will be performed on any untrusted ports within the VLAN. (Default: Disabled)

Web Interface

To configure global settings for DHCP Snooping:

1. Click IP Service, DHCP, Snooping.
2. Select Configure VLAN from the Step list.
3. Enable DHCP Snooping on any existing VLAN.
4. Click Apply

Figure 223: Configuring DHCP Snooping on a VLAN



IP Service > DHCP > Snooping

Step: 2. Configure VLAN

VLAN 1

DHCP Snooping Status Enabled

Apply Revert

Configuring Ports for DHCP Snooping Use the IP Service > DHCP > Snooping (Configure Interface) page to configure switch ports as trusted or untrusted.

Command Usage

- ◆ A trusted interface is an interface that is configured to receive only messages from within the network. An untrusted interface is an interface that is configured to receive messages from outside the network or fire wall.

- ◆ When DHCP snooping is enabled both globally and on a VLAN, DHCP packet filtering will be performed on any untrusted ports within the VLAN.
- ◆ When an untrusted port is changed to a trusted port, all the dynamic DHCP snooping bindings associated with this port are removed.
- ◆ Set all ports connected to DHCP servers within the local network or fire wall to trusted state. Set all other ports outside the local network or fire wall to untrusted state.

Parameters

These parameters are displayed:

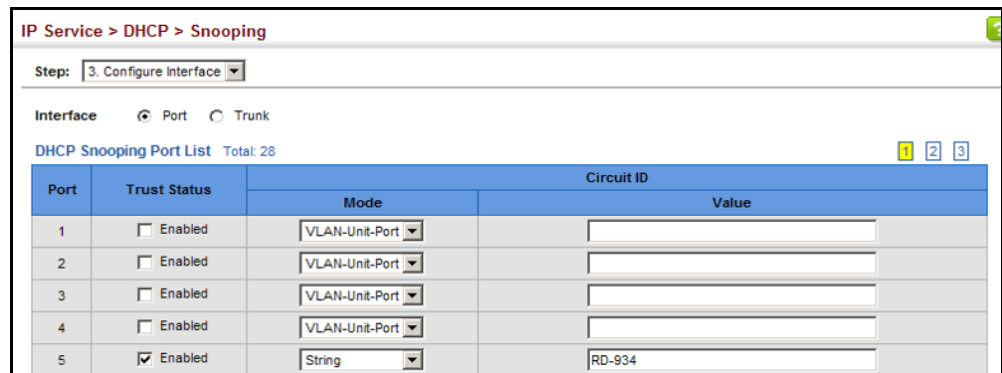
- ◆ **Trust Status** – Enables or disables a port as trusted. (Default: Disabled)
- ◆ **Circuit ID** – Specifies DHCP Option 82 circuit ID suboption information.
 - **Mode** – Specifies the default string “VLAN-Unit-Port” or an arbitrary string. (Default: VLAN-Unit-Port)
 - **Value** – An arbitrary string inserted into the circuit identifier field. (Range: 1-32 characters)

Web Interface

To configure global settings for DHCP Snooping:

1. Click IP Service, DHCP, Snooping.
2. Select Configure Interface from the Step list.
3. Set any ports within the local network or firewall to trusted.
4. Specify the mode used for sending circuit ID information, and an arbitrary string if required.
5. Click Apply

Figure 224: Configuring the Port Mode for DHCP Snooping



Displaying DHCP Snooping Binding Information

Use the IP Service > DHCP > Snooping (Show Information) page to display entries in the binding table.

Parameters

These parameters are displayed:

- ◆ **MAC Address** – Physical address associated with the entry.
- ◆ **IP Address** – IP address corresponding to the client.
- ◆ **Lease Time** – The time for which this IP address is leased to the client.
- ◆ **Type** – Entry types include:
 - **DHCP-Snooping** – Dynamically snooped.
 - **Static-DHCP-SNP** – Statically configured.
- ◆ **VLAN** – VLAN to which this entry is bound.
- ◆ **Interface** – Port or trunk to which this entry is bound.
- ◆ **Store** – Writes all dynamically learned snooping entries to flash memory. This function can be used to store the currently learned dynamic DHCP snooping entries to flash memory. These entries will be restored to the snooping table when the switch is reset. However, note that the lease time shown for a dynamic entry that has been restored from flash memory will no longer be valid.
- ◆ **Clear** – Removes all dynamically learned snooping entries from flash memory.

Web Interface

To display the binding table for DHCP Snooping:

1. Click IP Service, DHCP, Snooping.
2. Select Show Information from the Step list.
3. Use the Store or Clear function if required.

Figure 225: 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

Click the button to Store DHCP Snooping binding entries to flash.

Click the button to Clear DHCP Snooping binding entries from flash.

Basic Administration Protocols

This chapter describes basic administration tasks including:

- ◆ **Event Logging** – Sets conditions for logging event messages to system memory or flash memory, configures conditions for sending trap messages to remote log servers, and configures trap reporting to remote hosts using Simple Mail Transfer Protocol (SMTP).
- ◆ **Link Layer Discovery Protocol (LLDP)** – Configures advertisement of basic information about the local switch, or discovery of information about neighboring devices on the local broadcast domain.
- ◆ **Power over Ethernet⁶** – 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

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.

6. ECS4210-12P, ECS4210-28P

The System Logs page allows you to configure and limit system messages that are logged to flash or RAM memory. The default is for event levels 0 to 3 to be logged to flash and levels 0 to 7 to be logged to RAM.

Parameters

These parameters are displayed:

- ◆ **System Log Status** – Enables/disables the logging of debug or error messages to the logging process. (Default: Enabled)
- ◆ **Flash Level** – Limits log messages saved to the switch’s permanent flash memory for all levels up to the specified level. For example, if level 3 is specified, all messages from level 0 to level 3 will be logged to flash. (Range: 0-7, Default: 3)

Table 23: Logging Levels

Level	Severity Name	Description
7 ¹	Debug	Debugging messages
6	Informational	Informational messages only
5	Notice	Normal but significant condition, such as cold start
4	Warning	Warning conditions (e.g., return false, unexpected return)
3	Error	Error conditions (e.g., invalid input, default used)
2	Critical	Critical conditions (e.g., memory allocation, or free memory error - resource exhausted)
1	Alert	Immediate action needed
0	Emergency	System unusable

1. There are only Level 2, 5 and 6 error messages for the current firmware release.

- ◆ **RAM Level** – Limits log messages saved to the switch’s temporary RAM memory for all levels up to the specified level. For example, if level 7 is specified, all messages from level 0 to level 7 will be logged to RAM. (Range: 0-7, Default: 7)



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

Note: All log messages are retained in RAM and Flash after a warm restart (i.e., power is reset through the command interface).

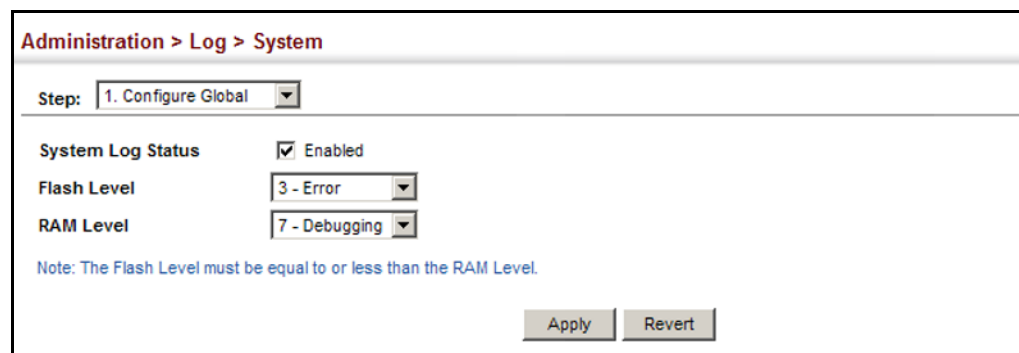
Note: All log messages are retained in Flash and purged from RAM after a cold restart (i.e., power is turned off and then on through the power source).

Web Interface

To configure the logging of error messages to system memory:

1. Click Administration, Log, System.
2. Select Configure Global from the Step list.
3. Enable or disable system logging, set the level of event messages to be logged to flash memory and RAM.
4. Click Apply.

Figure 226: Configuring Settings for System Memory Logs



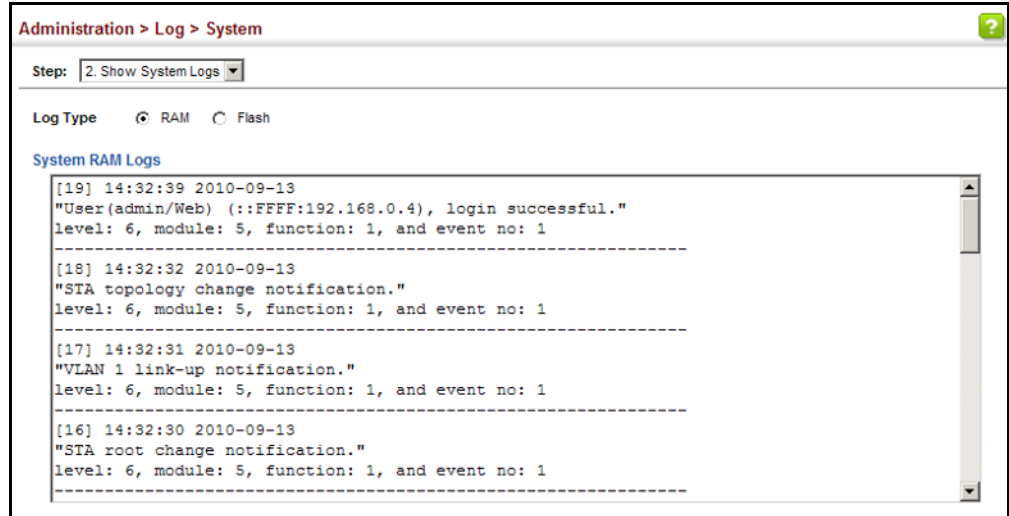
The screenshot shows a web interface for configuring system memory logs. The breadcrumb navigation is "Administration > Log > System". The "Step:" dropdown is set to "1. Configure Global". The "System Log Status" is checked and set to "Enabled". The "Flash Level" dropdown is set to "3 - Error". The "RAM Level" dropdown is set to "7 - Debugging". A note below the settings states: "Note: The Flash Level must be equal to or less than the RAM Level." At the bottom right, there are "Apply" and "Revert" buttons.

To show the error messages logged to system or flash memory:

1. Click Administration, Log, System.
2. Select Show 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 227: 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.

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 228: Configuring Settings for Remote Logging of Error Messages

Administration > Log > Remote

Remote Log Status	<input checked="" type="checkbox"/> Enabled		
Logging Facility	23 - Local use 7		
Logging Trap Level	7 - Debugging messages		
Server IP Address 1	192.168.0.4	Port	514
Server IP Address 2		Port	
Server IP Address 3		Port	
Server IP Address 4		Port	
Server IP Address 5		Port	

Apply Revert

Link Layer Discovery Protocol

Link Layer Discovery Protocol (LLDP) is used to discover basic information about neighboring devices on the local broadcast domain. LLDP is a Layer 2 protocol that uses periodic broadcasts to advertise information about the sending device. Advertised information is represented in Type Length Value (TLV) format according to the IEEE 802.1ab standard, and can include details such as device identification, capabilities and configuration settings. LLDP also defines how to store and maintain information gathered about the neighboring network nodes it discovers.

Link Layer Discovery Protocol - Media Endpoint Discovery (LLDP-MED) is an extension of LLDP intended for managing endpoint devices such as Voice over IP phones and network switches. The LLDP-MED TLVs advertise information such as network policy, power, inventory, and device location details. LLDP and LLDP-MED information can be used by SNMP applications to simplify troubleshooting, enhance network management, and maintain an accurate network topology.

Setting LLDP Timing Attributes

Use the Administration > LLDP (Configure Global) page to set attributes for general functions such as globally enabling LLDP on the switch, setting the message ageout time, and setting the frequency for broadcasting general advertisements or reports about changes in the LLDP MIB.

Parameters

These parameters are displayed:

- ◆ **LLDP** – Enables LLDP globally on the switch. (Default: Enabled)
- ◆ **Transmission Interval** – Configures the periodic transmit interval for LLDP advertisements. (Range: 5-32768 seconds; Default: 30 seconds)
- ◆ **Hold Time Multiplier** – Configures the time-to-live (TTL) value sent in LLDP advertisements as shown in the formula below. (Range: 2-10; Default: 4)

The time-to-live tells the receiving LLDP agent how long to retain all information pertaining to the sending LLDP agent if it does not transmit updates in a timely manner.

TTL in seconds is based on the following rule:

minimum value ((Transmission Interval * Holdtime Multiplier), or 65535)

Therefore, the default TTL is $4 * 30 = 120$ seconds.

- ◆ **Delay Interval** – Configures a delay between the successive transmission of advertisements initiated by a change in local LLDP MIB variables. (Range: 1-8192 seconds; Default: 2 seconds)

The transmit delay is used to prevent a series of successive LLDP transmissions during a short period of rapid changes in local LLDP MIB objects, and to increase the probability that multiple, rather than single changes, are reported in each transmission.

This attribute must comply with the rule:

$(4 * \text{Delay Interval}) \leq \text{Transmission Interval}$

- ◆ **Reinitialization Delay** – Configures the delay before attempting to re-initialize after LLDP ports are disabled or the link goes down. (Range: 1-10 seconds; Default: 2 seconds)

When LLDP is re-initialized on a port, all information in the remote systems LLDP MIB associated with this port is deleted.

- ◆ **Notification Interval** – Configures the allowed interval for sending SNMP notifications about LLDP MIB changes. (Range: 5-3600 seconds; Default: 5 seconds)

This parameter only applies to SNMP applications which use data stored in the LLDP MIB for network monitoring or management.

Information about changes in LLDP neighbors that occur between SNMP notifications is not transmitted. Only state changes that exist at the time of a notification are included in the transmission. An SNMP agent should therefore periodically check the value of `lldpStatsRemTableLastChangeTime` 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 229: Configuring LLDP Timing Attributes

The screenshot shows the 'Administration > LLDP' configuration page. At the top, there is a breadcrumb 'Administration > LLDP' and a 'Step:' dropdown menu set to '1. Configure Global'. Below this, the 'LLDP' section is checked and labeled 'Enabled'. Several timing parameters are listed with input fields and units:

Transmission Interval (5-32768)	30	sec
Hold Time Multiplier (2-10)	4	
Delay Interval (1-8192)	2	sec
Reinitialization Delay (1-10)	2	sec
Notification Interval (5-3600)	5	sec
MED Fast Start Count (1-10)	4	

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

Configuring LLDP Interface Attributes

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

Parameters

These parameters are displayed:

- ◆ **Admin Status** – Enables LLDP message transmit and receive modes for LLDP Protocol Data Units. (Options: Tx only, Rx only, TxRx, Disabled; Default: TxRx)

- ◆ **SNMP Notification** – Enables the transmission of SNMP trap notifications about LLDP and LLDP-MED changes. (Default: Disabled)

This option sends out SNMP trap notifications to designated target stations at the interval specified by the Notification Interval in the preceding section. Trap notifications include information about state changes in the LLDP MIB (IEEE 802.1AB), the LLDP-MED MIB (ANSI/TIA-1057), or vendor-specific LLDP-EXT-DOT1 and LLDP-EXT-DOT3 MIBs.

For information on defining SNMP trap destinations, see [“Specifying Trap Managers” on page 403](#).

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.

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.

- **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.
 - **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.
 - **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 66](#).
- ◆ **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 154](#)).
 - **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 135](#)).
 - **VLAN Name** – The name of all VLANs to which this interface has been assigned (see [“IEEE 802.1Q VLANs” on page 135](#)).
 - **Port and Protocol VLAN ID** – The port-based protocol VLANs configured on this interface (see [“IEEE 802.1Q VLANs” on page 135](#)).
- ◆ **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.
 - **Max Frame Size** – The maximum frame size. (See [“Configuring Support for Jumbo Frames” on page 68](#) for information on configuring the maximum frame size for this switch.
 - **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.
 - **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.

- ◆ **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.
 - **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).
 - **Inventory** – This option advertises device details useful for inventory management, such as manufacturer, model, software version and other pertinent information.
 - **Location** – This option advertises location identification details.
 - **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.
- ◆ **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 230: 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 24: LLDP MED Location CA Types

CA Type	Description	CA Value Example
1	National subdivisions (state, canton, province)	California
2	County, parish	Orange
3	City, township	Irvine
4	City division, borough, city district	West Irvine
5	Neighborhood, block	Riverside
6	Group of streets below the neighborhood level	Exchange

Table 24: LLDP MED Location CA Types (Continued)

CA Type	Description	CA Value Example
18	Street suffix or type	Avenue
19	House number	320
20	House number suffix	A
21	Landmark or vanity address	Tech Center
26	Unit (apartment, suite)	Apt 519
27	Floor	5
28	Room	509B

- ◆ Any number of CA type and value pairs can be specified for the civic address location, as long as the total does not exceed 250 characters.

Parameters

These parameters are displayed:

- ◆ **CA-Type** – Descriptor of the data civic address value. (Range: 0-255)
- ◆ **CA-Value** – Description of a location. (Range: 1-32 characters)

Web Interface

To specify the physical location of the attached device:

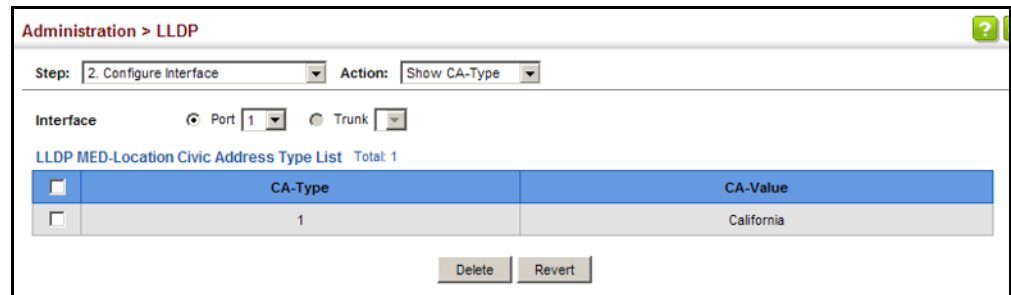
1. Click Administration, LLDP.
2. Select Configure Interface from the Step list.
3. Select Add CA-Type from the Action list.
4. Select an interface from the Port or Trunk list.
5. Specify a CA-Type and CA-Value pair.
6. Click Apply.

Figure 231: Configuring the Civic Address for an LLDP Interface

To show the physical location of the attached device:

1. Click Administration, LLDP.
2. Select Configure Interface from the Step list.
3. Select Show CA-Type from the Action list.
4. Select an interface from the Port or Trunk list.

Figure 232: 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:

Global Settings

- ◆ **Chassis Type** – Identifies the chassis containing the IEEE 802 LAN entity associated with the transmitting LLDP agent. There are several ways in which a chassis may be identified and a chassis ID subtype is used to indicate the type of component being referenced by the chassis ID field.

Table 25: Chassis ID Subtype

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

- ◆ **Chassis ID** – An octet string indicating the specific identifier for the particular chassis in this system.
- ◆ **System Name** – A string that indicates the system’s administratively assigned name (see “[Displaying System Information](#)” on page 66).
- ◆ **System Description** – A textual description of the network entity. This field is also displayed by the **show system** command.
- ◆ **System Capabilities Supported** – The capabilities that define the primary function(s) of the system.

Table 26: System Capabilities

ID Basis	Reference
Other	—
Repeater	IETF RFC 2108
Bridge	IETF RFC 2674
WLAN Access Point	IEEE 802.11 MIB
Router	IETF RFC 1812
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 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.

Interface Settings

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

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

Interface Details

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

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

Table 27: Port ID Subtype

ID Basis	Reference
Interface alias	IfAlias (IETF RFC 2863)
Port component	EntPhysicalAlias when entPhysicalClass has a value 'port(10)' or 'backplane(4)' (IETF RFC 4133)
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, or Trunk.

Figure 233: Displaying Local Device Information for LLDP (General)

Administration > LLDP

Step: 3. Show Local Device Information

General
 Port
 Port Details
 Trunk
 Trunk Details

LLDP Local Device Information

Chassis Type	MAC Address
Chassis ID	00-E0-0C-00-00-FD
System Name	
System Description	Managed 24G+4GSFP Switch
System Capabilities Supported	Bridge
System Capabilities Enabled	Bridge
Management Address	192.168.0.2 (IPv4)

Figure 234: 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	70-72-CF-5B-DA-D5
2	Ethernet Port on unit 1, port 2	70-72-CF-5B-DA-D6
3	Ethernet Port on unit 1, port 3	70-72-CF-5B-DA-D7
4	Ethernet Port on unit 1, port 4	70-72-CF-5B-DA-D8
5	Ethernet Port on unit 1, port 5	70-72-CF-5B-DA-D9

Figure 235: 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-E0-0C-00-01-FE
Port Description	Ethernet Port on unit 1, port 1
MED Capability	LLDP-MED Capabilities, Network Policy, Location Identification, Extended Power via MDI - PSE, Extended Power via MDI - PD, Inventory

Displaying LLDP Remote Device Information Use the Administration > LLDP (Show Remote Device Information) page to display information about devices connected directly to the switch's ports which are advertising information through LLDP, or to display detailed information about an LLDP-enabled device connected to a specific port on the local switch.

Parameters

These parameters are displayed:

Port

- ◆ **Local Port** – The local port to which a remote LLDP-capable device is attached.
- ◆ **Chassis ID** – An octet string indicating the specific identifier for the particular chassis in this system.
- ◆ **Port ID** – A string that contains the specific identifier for the port from which this LLDPDU was transmitted.
- ◆ **System Name** – A string that indicates the system's administratively assigned name.

Port Details

- ◆ **Port** – Port identifier on local switch.
- ◆ **Remote Index** – Index of remote device attached to this port.
- ◆ **Local Port** – The local port to which a remote LLDP-capable device is attached.
- ◆ **Chassis Type** – Identifies the chassis containing the IEEE 802 LAN entity associated with the transmitting LLDP agent. There are several ways in which a chassis may be identified and a chassis ID subtype is used to indicate the type of component being referenced by the chassis ID field. (See [Table 25, "Chassis ID Subtype," on page 367.](#))
- ◆ **Chassis ID** – An octet string indicating the specific identifier for the particular chassis in this system.
- ◆ **System Name** – A string that indicates the system's assigned name.
- ◆ **System Description** – A textual description of the network entity.
- ◆ **Port Type** – Indicates the basis for the identifier that is listed in the Port ID field. See [Table 27, "Port ID Subtype," on page 369.](#)
- ◆ **Port Description** – A string that indicates the port's description. If RFC 2863 is implemented, the ifDescr object should be used for this field.
- ◆ **Port ID** – A string that contains the specific identifier for the port from which this LLDPDU was transmitted.

- ◆ **System Capabilities Supported** – The capabilities that define the primary function(s) of the system. (See Table 26, "System Capabilities," on page 368.)
- ◆ **System Capabilities Enabled** – The primary function(s) of the system which are currently enabled. (See Table 26, "System Capabilities," on page 368.)
- ◆ **Management Address List** – The management addresses for this device. Since there are typically a number of different addresses associated with a Layer 3 device, an individual LLDP PDU may contain more than one management address TLV.

If no management address is available, the address should be the MAC address for the CPU or for the port sending this advertisement.

Port Details – 802.1 Extension Information

- ◆ **Remote Port VID** – The port's default VLAN identifier (PVID) indicates the VLAN with which untagged or priority-tagged frames are associated.
- ◆ **Remote Port-Protocol VLAN List** – The port-based protocol VLANs configured on this interface, whether the given port (associated with the remote system) supports port-based protocol VLANs, and whether the port-based protocol VLANs are enabled on the given port associated with the remote system.
- ◆ **Remote VLAN Name List** – VLAN names associated with a port.
- ◆ **Remote Protocol Identity List** – Information about particular protocols that are accessible through a port. This object represents an arbitrary local integer value used by this agent to identify a particular protocol identity, and an octet string used to identify the protocols associated with a port of the remote system.

Port Details – 802.3 Extension Port Information

- ◆ **Remote Port Auto-Neg Supported** – Shows whether the given port (associated with remote system) supports auto-negotiation.
- ◆ **Remote Port Auto-Neg Adv-Capability** – The value (bitmap) of the ifMauAutoNegCapAdvertisedBits object (defined in IETF RFC 3636) which is associated with a port on the remote system.

Table 28: Remote Port Auto-Negotiation Advertised Capability

Bit	Capability
0	other or unknown
1	10BASE-T half duplex mode
2	10BASE-T full duplex mode
3	100BASE-T4
4	100BASE-TX half duplex mode

Table 28: 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 Controlable** – 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 Capability⁷

- ◆ **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.

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

Port Details – Network Policy⁷

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

Port Details – Location Identification⁷

- ◆ **Location Data Format** – Any of these location ID data formats:
 - Coordinate-based LCI⁸ – Defined in RFC 3825, includes latitude resolution, latitude, longitude resolution, longitude, altitude type, altitude resolution, altitude, and datum.
 - Civic Address LCI⁸ – Includes What, Country code, CA type, CA length and CA value. “What” is described as the field entry “Device entry refers to” under “[Configuring LLDP Interface Attributes](#).” The other items are described under “[Configuring LLDP Interface Civic-Address](#).”

8. 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 – Inventory⁷

- ◆ **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 basic LLDP information for a remote interface:

1. Click Administration, LLDP.
2. Select Show Remote Device Information from the Step list.
3. Select Port or Trunk.

Figure 236: Displaying Basic LLDP Information for a Remote Device (Port)

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

To display detailed LLDP information for a remote Interface:

- 1.** Click Administration, LLDP.
- 2.** Select Show Remote Device Information from the Step list.
- 3.** Select Port Details or Trunk Details.
- 4.** Select a port on this switch and the index for a remote device attached to this interface.
- 5.** Click Query.

Figure 237: Displaying Remote Device Information for LLDP (Port Details)

Administration > LLDP

Step: 4. Show Remote Device Information

Port
 Port Details
 Trunk
 Trunk Details

Port: 1

Remote Index: 6

Query

LLDP Remote Device Port Information

Local Port	1	Port Type	MAC Address
Chassis Type	MAC Address	Port Description	Ethernet Port on unit 1, port 1
Chassis ID	B4-0E-DC-39-F4-4B	Port ID	B4-0E-DC-39-F4-4C
System Name		System Capabilities Supported	Bridge
System Description		System Capabilities Enabled	Bridge

Management Address List Total: 1

Address	Address Type
192.168.0.3	IPv4 Address

802.1 Extension Information

Remote Port VID: 1

Remote Port-Protocol VLAN List Total: 1

VLAN	Support	Status
3	Yes	Enabled

Remote VLAN Name List Total: 3

VLAN	Name
1	DefaultVlan
2	R&D
3	Protocol

Remote Protocol Identity List Total: 1

Remote Protocol Identity
88-CC

802.3 Extension Port Information

Remote Port Auto-Neg Supported	Yes	Remote Port Auto-Neg Status	Enabled
Remote Port Auto-Neg Adv-Capability	0000	Remote Port MAU Type	6

802.3 Extension Power Information

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

802.3 Extension Trunk Information

Remote Link Aggregation Capable	Yes	Remote Link Aggregation Status	Disabled
Remote Link Port ID	0		

802.3 Extension Frame Information

Remote Max Frame Size	1518
-----------------------	------

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

Figure 238: Displaying Remote Device Information for LLDP (End Node)

Administration > LLDP			
Step: 4. Show Remote Device Information			
LLDP-MED Capability			
Device Class	Network Connectivity		
Supported Capabilities	LLDP-MED Capabilities, Network Policy, Location Identification, Inventory		
Current Capabilities	LLDP-MED Capabilities, Network Policy, Location Identification, Inventory		
Network Policy			
Application Type	Guest Voice Signaling	Unknown Policy Flag	Disabled
Tagged Flag	Disabled	VLAN ID	7
Layer 2 Priority	2	DSCP Value	62
Location Identification			
Location Data Format	Coordinate-based LCI		
Country Code	TW	What	2
Inventory			
Hardware Revision	R01	Firmware Revision	1.0.0.2
Software Revision	1.0.0.2	Serial Number	LN11360051
Manufacture Name		Model Name	
Asset ID	1		

Displaying Device Statistics

Use the Administration > LLDP (Show Device Statistics) page to display statistics for LLDP-capable devices attached to the switch, and for LLDP protocol messages transmitted or received on all local interfaces.

Parameters

These parameters are displayed:

General Statistics on Remote Devices

- ◆ **Neighbor Entries List Last Updated** – The time the LLDP neighbor entry list was last updated.
- ◆ **New Neighbor Entries Count** – The number of LLDP neighbors for which the remote TTL has not yet expired.
- ◆ **Neighbor Entries Deleted Count** – The number of LLDP neighbors which have been removed from the LLDP remote systems MIB for any reason.
- ◆ **Neighbor Entries Dropped Count** – The number of times which the remote database on this switch dropped an LLDPDU because of insufficient resources.
- ◆ **Neighbor Entries Age-out Count** – The number of times that a neighbor's information has been deleted from the LLDP remote systems MIB because the remote TTL timer has expired.

Port/Trunk

- ◆ **Frames Discarded** – Number of frames discarded because they did not conform to the general validation rules as well as any specific usage rules defined for the particular TLV.
- ◆ **Frames Invalid** – A count of all LLDPDUs received with one or more detectable errors.
- ◆ **Frames Received** – Number of LLDP PDUs received.
- ◆ **Frames Sent** – Number of LLDP PDUs transmitted.
- ◆ **TLVs Unrecognized** – A count of all TLVs not recognized by the receiving LLDP local agent.
- ◆ **TLVs Discarded** – A count of all LLDPDUs received and then discarded due to insufficient memory space, missing or out-of-sequence attributes, or any other reason.
- ◆ **Neighbor Ageouts** – A count of the times that a neighbor's information has been deleted from the LLDP remote systems MIB because the remote TTL timer has expired.

Web Interface

To display statistics for LLDP-capable devices attached to the switch:

1. Click Administration, LLDP.
2. Select Show Device Statistics from the Step list.
3. Select General, Port, or Trunk.

Figure 239: Displaying LLDP Device Statistics (General)

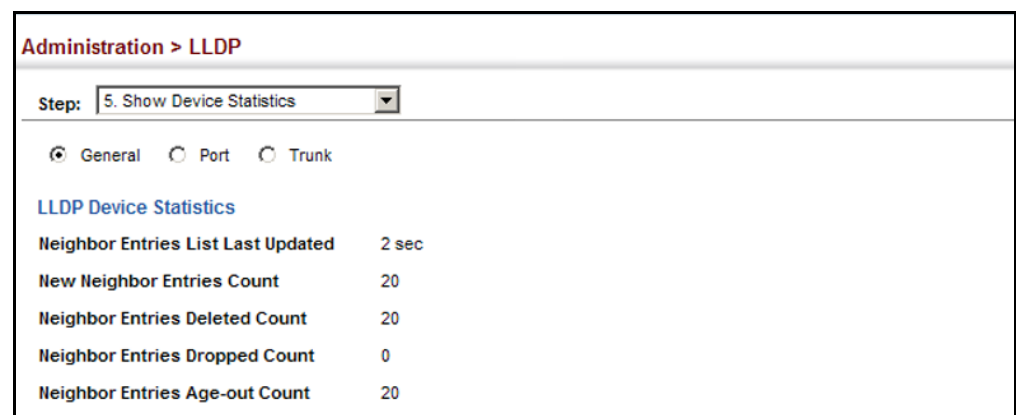
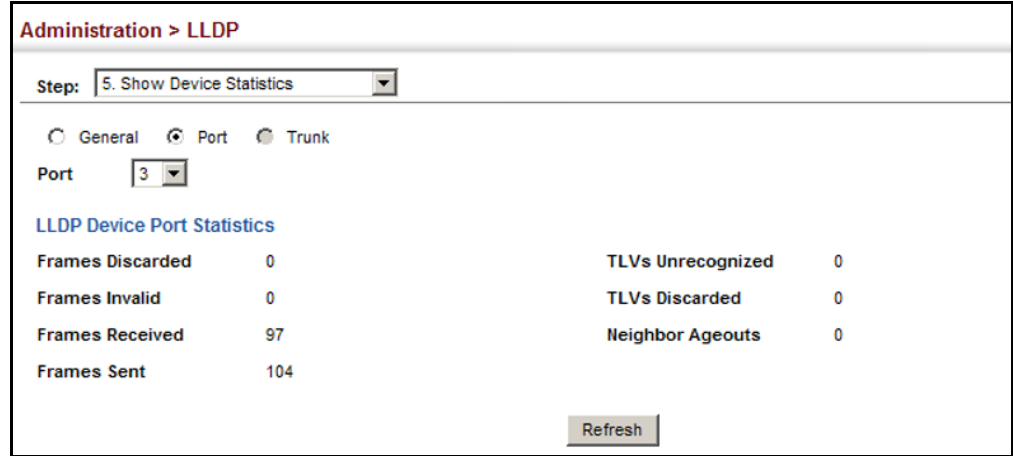


Figure 240: Displaying LLDP Device Statistics (Port)



Power over Ethernet

The ECS4210-12P and ECS4210-28P 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 (IEEE 802.3af or 802.3at).

The switch's power management allows 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.



Note: For more information on using the PoE provided by this switch refer to the *Installation Guide*.

Displaying the Switch's Overall PoE Power Budget

Use the Administration > PoE (Configure Global) page to display 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:

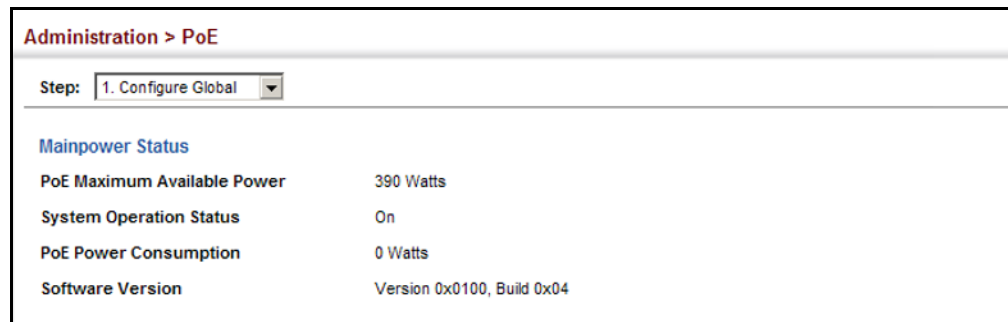
- ◆ **PoE Maximum Available Power** – The power budget for the switch. If devices connected to the switch require more power than the switch budget, the port power priority settings are used to control the supplied power.
(Fixed: 390 Watts)
- ◆ **System Operation Status** – Status of the PoE power service 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.

Figure 241: Showing the Switch's PoE Budget



Setting The Port PoE Power Budget Use the Administration > PoE (Configure Interface) page to set the 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.
- ◆ All the RJ-45 ports support both the IEEE 802.3af and IEEE 802.3at standards. The total PoE power delivered by all ports cannot exceed the maximum power budget of 390W. This means that up to 11 ports can supply a maximum 34.2W of power simultaneously to connected devices (802.3at), or all 12/24 ports can supply up to 15.4W (802.3af).
- ◆ If the power demand from devices connected to all switch ports exceeds the power budget as determined during bootup, 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, power is provided to the port only if the switch can drop power to one or more lower-priority ports and thereby remain within its overall budget.

Power will be dropped from low-priority ports in sequence starting from port number 1.

If priority is not set for any ports, power is denied in reverse sequence, starting from Port 12/24.
 - If sufficient power cannot be freed up for a critical or high-priority port by turning off power to lower-priority ports, power will not be supplied to the newly connected device.
- ◆ If priority is not set for any ports, and there is not sufficient power to supply all of the ports, port priority defaults to Port 1, Port 2, Port 3 ... Port 24, with available power being supplied in that sequence.
- ◆ If a device is connected to a switch port after bootup and the switch detects that it requires more than the power budget set for the port or for the overall switch, no power is supplied to the device regardless of its priority setting.

Parameters

These parameters are displayed:

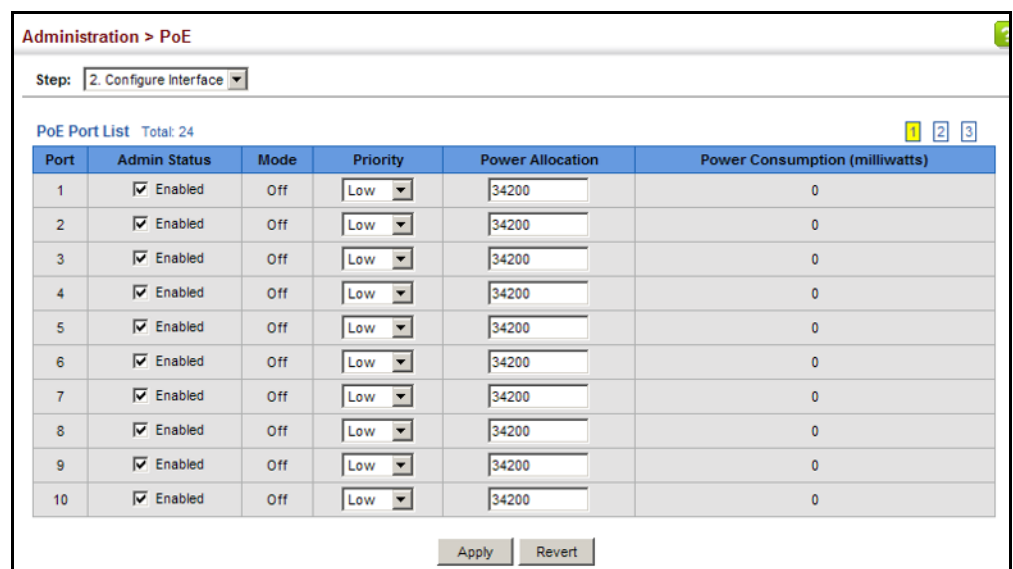
- ◆ **Port** – The port number on the switch.
- ◆ **Admin Status** – Enables PoE power on a port. Power is automatically supplied when a device is detected on a port, providing that the power demanded does not exceed the switch or port power budget. (Default: Enabled)
- ◆ **Mode** – Shows whether or not PoE power is being supplied to a port.
- ◆ **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: 3000-34200 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.
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 242: Setting a Port's PoE Budget



The screenshot shows a web interface for configuring PoE settings. The breadcrumb is "Administration > PoE". The current step is "2. Configure Interface". Below this is a "PoE Port List" with a total of 24 ports. The table has columns for Port, Admin Status, Mode, Priority, Power Allocation, and Power Consumption (milliwatts). The table shows 10 rows of ports, all with Admin Status checked and Enabled, Mode set to Off, Priority set to Low, Power Allocation set to 34200, and Power Consumption set to 0. At the bottom of the table are "Apply" and "Revert" buttons.

Port	Admin Status	Mode	Priority	Power Allocation	Power Consumption (milliwatts)
1	<input checked="" type="checkbox"/> Enabled	Off	Low	34200	0
2	<input checked="" type="checkbox"/> Enabled	Off	Low	34200	0
3	<input checked="" type="checkbox"/> Enabled	Off	Low	34200	0
4	<input checked="" type="checkbox"/> Enabled	Off	Low	34200	0
5	<input checked="" type="checkbox"/> Enabled	Off	Low	34200	0
6	<input checked="" type="checkbox"/> Enabled	Off	Low	34200	0
7	<input checked="" type="checkbox"/> Enabled	Off	Low	34200	0
8	<input checked="" type="checkbox"/> Enabled	Off	Low	34200	0
9	<input checked="" type="checkbox"/> Enabled	Off	Low	34200	0
10	<input checked="" type="checkbox"/> Enabled	Off	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 29: SNMPv3 Security Models and Levels

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



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

Command Usage

Configuring SNMPv1/2c Management Access

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

1. Use the Administration > SNMP (Configure Global) page to enable SNMP on the switch, and to enable trap messages.
2. Use the Administration > SNMP (Configure User - Add Community) page to configure the community strings authorized for management access.
3. Use the Administration > SNMP (Configure Trap) page to specify trap managers so that key events are reported by this switch to your management station.

Configuring SNMPv3 Management Access

1. Use the Administration > SNMP (Configure Global) page to enable SNMP on the switch, and to enable trap messages.
2. Use the Administration > SNMP (Configure Trap) page to specify trap managers so that key events are reported by this switch to your management station.

3. Use the Administration > SNMP (Configure Engine) page to change the local engine ID. If you want to change the default engine ID, it must be changed before configuring other parameters.
4. Use the Administration > SNMP (Configure View) page to specify read and write access views for the switch MIB tree.
5. Use the Administration > SNMP (Configure User) page to configure SNMP user groups with the required security model (i.e., SNMP v1, v2c or v3) and security level (i.e., authentication and privacy).
6. Use the Administration > SNMP (Configure Group) page to assign SNMP users to groups, along with their specific authentication and privacy passwords.

Configuring Global Settings for SNMP

Use the Administration > SNMP (Configure Global) page to enable SNMPv3 service for all management clients (i.e., versions 1, 2c, 3), and to enable trap messages.

Parameters

These parameters are displayed:

- ◆ **Agent Status** – Enables SNMP on the switch. (Default: Enabled)
- ◆ **Authentication Traps⁹** – Issues a notification message to specified IP trap managers whenever an invalid community string is submitted during the SNMP access authentication process. (Default: Enabled)
- ◆ **Link-up and Link-down Traps⁹** – Issues a notification message whenever a port link is established or broken. (Default: Enabled)

Web Interface

To configure global settings for SNMP:

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

9. 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 390).

Figure 243: Configuring Global Settings for SNMP

Administration > SNMP

Step: 1. Configure Global

Agent Status	<input checked="" type="checkbox"/> Enabled
Authentication Traps	<input checked="" type="checkbox"/> Enabled
Link-up and Link-down Traps	<input checked="" type="checkbox"/> Enabled

Apply Revert

Setting the Local Engine ID

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

Command Usage

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

Parameters

These parameters are displayed:

- ◆ **Engine ID** – A new engine ID can be specified by entering 9 to 64 hexadecimal characters (5 to 32 octets in hexadecimal format). If an odd number of characters are specified, a trailing zero is added to the value to fill in the last octet. For example, the value “123456789” is equivalent to “1234567890”.
- ◆ **Engine Boots** – The number of times that the engine has (re-)initialized since the SNMP EngineID 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 244: Configuring the Local Engine ID for SNMP

The screenshot shows a web interface for configuring SNMP. At the top, it says 'Administration > SNMP'. Below that, there are two dropdown menus: 'Step: 2. Configure Engine' and 'Action: Set Engine ID'. The main area contains two fields: 'Engine ID' with the value '800001030300000c0000fd0000' and 'Engine Boots' with the value '5'. At the bottom right, there are two buttons: 'Default' and 'Save'.

Specifying a Remote Engine ID

Use the Administration > SNMP (Configure Engine - Add Remote Engine) page to configure a engine ID for a remote management station. To allow management access from an SNMPv3 user on a remote device, you must first specify the engine identifier for the SNMP agent on the remote device where the user resides. The remote engine ID is used to compute the security digest for authentication and encryption of packets passed between the switch and a user on the remote host.

Command Usage

- ◆ SNMP passwords are localized using the engine ID of the authoritative agent. For informs, the authoritative SNMP agent is the remote agent. You therefore need to configure the remote agent’s SNMP engine ID before you can send proxy requests or informs to it. (See [“Configuring Remote SNMPv3 Users” on page 401.](#))

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 IP 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 245: Configuring a Remote Engine ID for SNMP

Administration > SNMP

Step: 2. Configure Engine Action: Add Remote Engine

Remote Engine ID: 5432100000

Remote IP Host: 192.168.1.19

Apply Revert

To show the remote SNMP engine IDs:

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

Figure 246: Showing Remote Engine IDs for SNMP

Administration > SNMP

Step: 2. Configure Engine Action: Show Remote Engine

SNMPv3 Remote Engine List Total: 1

<input type="checkbox"/>	Remote Engine ID	Remote IP Host
<input type="checkbox"/>	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. (Range: 1-64 characters)
- ◆ **OID Subtree** – Specifies the initial object identifier of a branch within the MIB tree. Wild cards can be used to mask a specific portion of the OID string. Use the Add OID Subtree page to configure additional object identifiers. (Range: 1-64 characters).
- ◆ **Type** – Indicates if the object identifier of a branch within the MIB tree is included or excluded from the SNMP view.

Add OID Subtree

- ◆ **View Name** – Lists the SNMP views configured in the Add View page. (Range: 1-32 characters).
- ◆ **OID Subtree** – Adds an additional object identifier of a branch within the MIB tree to the selected View. Wild cards can be used to mask a specific portion of the OID string. (Range: 1-64 characters).
- ◆ **Type** – Indicates if the object identifier of a branch within the MIB tree is included or excluded from the SNMP view.

Web Interface

To configure an SNMP view of the switch's MIB database:

1. Click Administration, SNMP.
2. Select Configure View from the Step list.
3. Select Add View from the Action list.
4. Enter a view name and specify the initial OID subtree in the switch's MIB database to be included or excluded in the view. Use the Add OID Subtree page to add additional object identifier branches to the view.
5. Click Apply

Figure 247: Creating an SNMP View

Administration > SNMP

Step: 3. Configure View Action: Add View

View Name: ifEntry.a

OID Subtree: 1.3.6.1.2.1.2.2.1.1.*

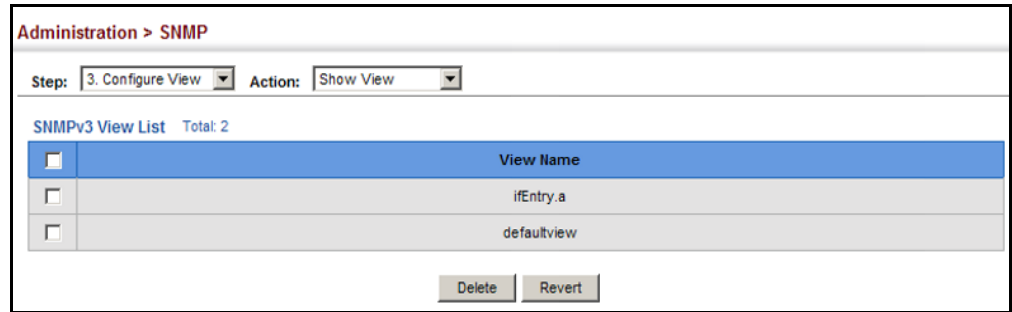
Type: Included

Apply Revert

To show the SNMP views of the switch's MIB database:

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

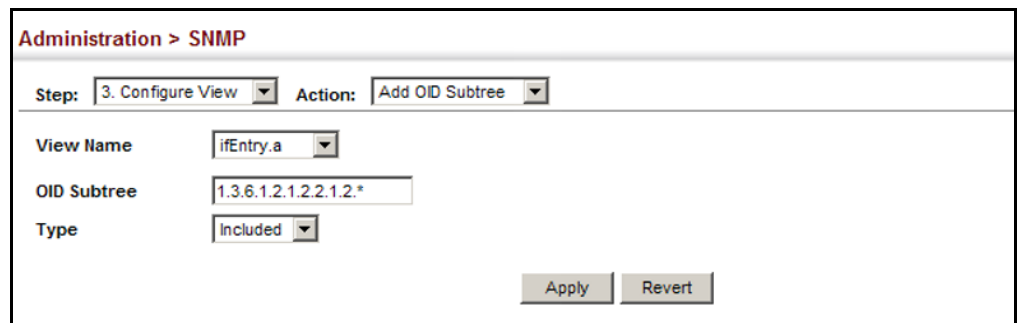
Figure 248: Showing SNMP Views



To add an object identifier to an existing SNMP view of the switch's MIB database:

1. Click Administration, SNMP.
2. Select Configure View from the Step list.
3. Select Add OID Subtree from the Action list.
4. Select a view name from the list of existing views, and specify an additional OID subtree in the switch's MIB database to be included or excluded in the view.
5. Click Apply

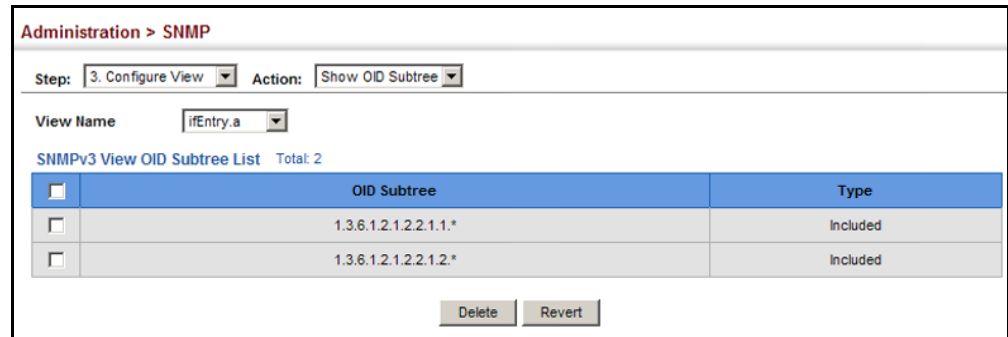
Figure 249: Adding an OID Subtree to an SNMP View



To show the OID branches configured for the SNMP views of the switch's MIB database:

1. Click Administration, SNMP.
2. Select Configure View from the Step list.
3. Select Show OID Subtree from the Action list.
4. Select a view name from the list of existing views.

Figure 250: Showing the OID Subtree Configured for SNMP Views



Configuring SNMPv3 Groups

Use the Administration > SNMP (Configure Group) page to add an SNMPv3 group which can be used to set the access policy for its assigned users, restricting them to specific read, write, and notify views. You can use the pre-defined default groups or create new groups to map a set of SNMP users to SNMP views.

Parameters

These parameters are displayed:

- ◆ **Group Name** – The name of the SNMP group to which the user is assigned. (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: 1-64 characters)
- ◆ **Write View** – The configured view for write access. (Range: 1-64 characters)
- ◆ **Notify View** – The configured view for notifications. (Range: 1-64 characters)

Table 30: Supported Notification Messages

Model	Level	Group
<i>RFC 1493 Traps</i>		
newRoot	1.3.6.1.2.1.17.0.1	The newRoot trap indicates that the sending agent has become the new root of the Spanning Tree; the trap is sent by a bridge soon after its election as the new root, e.g., upon expiration of the Topology Change Timer immediately subsequent to its election.
topologyChange	1.3.6.1.2.1.17.0.2	A topologyChange trap is sent by a bridge when any of its configured ports transitions from the Learning state to the Forwarding state, or from the Forwarding state to the Discarding state. The trap is not sent if a newRoot trap is sent for the same transition.
<i>SNMPv2 Traps</i>		
coldStart	1.3.6.1.6.3.1.1.5.1	A coldStart trap signifies that the SNMPv2 entity, acting in an agent role, is reinitializing itself and that its configuration may have been altered.
warmStart	1.3.6.1.6.3.1.1.5.2	A warmStart trap signifies that the SNMPv2 entity, acting in an agent role, is reinitializing itself such that its configuration is unaltered.
linkDown ¹	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	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	1.3.6.1.6.3.1.1.5.5	An authenticationFailure trap signifies that the SNMPv2 entity, acting in an agent role, has received a protocol message that is not properly authenticated. While all implementations of the SNMPv2 must be capable of generating this trap, the snmpEnableAuthenTraps object indicates whether this trap will be generated.
<i>RMON Events (V2)</i>		
risingAlarm	1.3.6.1.2.1.16.0.1	The SNMP trap that is generated when an alarm entry crosses its rising threshold and generates an event that is configured for sending SNMP traps.
fallingAlarm	1.3.6.1.2.1.16.0.2	The SNMP trap that is generated when an alarm entry crosses its falling threshold and generates an event that is configured for sending SNMP traps.

Table 30: Supported Notification Messages (Continued)

Model	Level	Group
<i>Private Traps²</i>		
swPowerStatusChangeTrap	1.3.6.1.4.1.259.6.10.112.2.1.0.1	This trap is sent when the power state changes.
swPortSecurityTrap	1.3.6.1.4.1.259.10.1.112.2.1.0.36	This trap is sent when the port is being intruded. This trap will only be sent when the portSecActionTrap is enabled.
swIpFilterRejectTrap	1.3.6.1.4.1.259.6.10.112.2.1.0.40	This trap is sent when an incorrect IP address is rejected by the IP filter.
swAuthenticationFailure	1.3.6.1.4.1.259.10.1.112.2.1.0.66	This trap will be triggered if authentication fails.
swAuthenticationSuccess	1.3.6.1.4.1.259.10.1.112.2.1.0.67	This trap will be triggered if authentication is successful.
swAtcBcastStormAlarmFireTrap	1.3.6.1.4.1.259.10.1.112.2.1.0.70	When broadcast traffic is detected as a storm, this trap is fired.
swAtcBcastStormAlarmClearTrap	1.3.6.1.4.1.259.10.1.112.2.1.0.71	When a broadcast storm is detected as normal traffic, this trap is fired.
swAtcBcastStormTcApplyTrap	1.3.6.1.4.1.259.10.1.112.2.1.0.72	When ATC is activated, this trap is fired.
swAtcBcastStormTcReleaseTrap	1.3.6.1.4.1.259.10.1.112.2.1.0.73	When ATC is released, this trap is fired.
swAtcMcastStormAlarmFireTrap	1.3.6.1.4.1.259.10.1.112.2.1.0.74	When multicast traffic is detected as the storm, this trap is fired.
swAtcMcastStormAlarmClearTrap	1.3.6.1.4.1.259.10.1.112.2.1.0.75	When multicast storm is detected as normal traffic, this trap is fired.
swAtcMcastStormTcApplyTrap	1.3.6.1.4.1.259.10.1.112.2.1.0.76	When ATC is activated, this trap is fired.
swAtcMcastStormTcReleaseTrap	1.3.6.1.4.1.259.10.1.112.2.1.0.77	When ATC is released, this trap is fired.
swLoopbackDetectionTrap	1.3.6.1.4.1.259.10.1.112.2.1.0.95	This trap is sent when loopback BPDUs have been detected.
networkAccessPortLinkDetectionTrap	1.3.6.1.4.1.259.6.10.112.2.1.0.96	This trap is sent when a networkAccessPortLinkDetection event is triggered.
autoUpgradeTrap	1.3.6.1.4.1.259.10.1.112.2.1.0.104	This trap is sent when auto upgrade is executed.
swCpuUtiRisingNotification	1.3.6.1.4.1.259.10.1.112.2.1.0.107	This notification indicates that the CPU utilization has risen from cpuUtiFallingThreshold to cpuUtiRisingThreshold.
swCpuUtiFallingNotification	1.3.6.1.4.1.259.10.1.112.2.1.0.108	This notification indicates that the CPU utilization has fallen from cpuUtiRisingThreshold to cpuUtiFallingThreshold.
swMemoryUtiRisingThresholdNotification	1.3.6.1.4.1.259.10.1.112.2.1.0.109	This notification indicates that the memory utilization has risen from memoryUtiFallingThreshold to memoryUtiRisingThreshold.
swMemoryUtiFallingThresholdNotification	1.3.6.1.4.1.259.10.1.112.2.1.0.110	This notification indicates that the memory utilization has fallen from memoryUtiRisingThreshold to memoryUtiFallingThreshold.

Table 30: Supported Notification Messages (Continued)

Model	Level	Group
dhcpRogueServerAttackTrap	1.3.6.1.4.1.259.6.10.112.2.1.0.114	This trap is sent when receiving a DHCP packet from a rogue server.
swLoginFailureTrap	1.3.6.1.4.1.259.6.10.112.2.1.0.139	This trap is sent when login fail via console, telnet, or web.
swLoginSucceedTrap	1.3.6.1.4.1.259.6.10.112.2.1.0.140	his trap is sent when login succeed via console, telnet, or web.
lbdDetectionTrap	1.3.6.1.4.1.259.6.10.112.2.1.0.141	This trap is sent when a loopback condition is detected by LBD.
lbdRecoveryTrap	1.3.6.1.4.1.259.6.10.112.2.1.0.142	This trap is sent when a recovery is done by LBD.
udldPortShutdownTrap	1.3.6.1.4.1.259.6.10.112.2.1.0.192	This trap is sent when the port is shut down by UDLD.
swAmtrHwSecurityIntrusionMacTrap	1.3.6.1.4.1.259.6.10.112.2.1.0.204	This trap is sent when learned MAC addresses exceed hardware security MAC max count.
swAmtrHwSecurityMacPortMovedTrap	1.3.6.1.4.1.259.6.10.112.2.1.0.205	This trap is sent when a hardware security static/dynamic MAC address has been moved to another port.

1. These are legacy notifications and therefore must be enabled in conjunction with the corresponding traps on the SNMP Configuration menu.

2. The MIB OID for
 ECS4210-12P is 1.3.6.1.4.1.259.10.1.42.104
 ECS4210-12T is 1.3.6.1.4.1.259.10.1.42.103
 ECS4210-28T is 1.3.6.1.4.1.259.10.1.42.101
 ECS4210-28P is 1.3.6.1.4.1.259.10.1.42.102

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 251: Creating an SNMP Group

To show SNMP groups:

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

Figure 252: Showing SNMP Groups

<input type="checkbox"/>	Group Name	Model	Level	Read View	Write View	Notify View
<input type="checkbox"/>	public	v1	noAuthNoPriv	defaultview	none	none
<input type="checkbox"/>	public	v2c	noAuthNoPriv	defaultview	none	none
<input type="checkbox"/>	private	v1	noAuthNoPriv	defaultview	defaultview	none
<input type="checkbox"/>	private	v2c	noAuthNoPriv	defaultview	defaultview	none
<input type="checkbox"/>	secure-users	v3	authPriv	ifEntry.a	ifEntry.a	ifEntry.a

Setting Community Access Strings Use the Administration > SNMP (Configure User - Add Community) page to configure up to five community strings authorized for management access by clients using SNMP v1 and v2c. For security reasons, you should consider removing the default strings.

Parameters

These parameters are displayed:

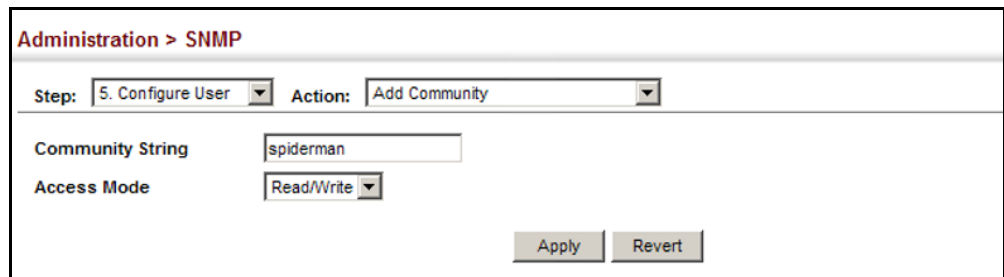
- ◆ **Community String** – A community string that acts like a password and permits access to the SNMP protocol.
Range: 1-32 characters, case sensitive
Default strings: “public” (Read-Only), “private” (Read/Write)
- ◆ **Access Mode** – Specifies the access rights for the community string:
 - **Read-Only** – Authorized management stations are only able to retrieve MIB objects.
 - **Read/Write** – Authorized management stations are able to both retrieve and modify MIB objects.

Web Interface

To set a community access string:

1. Click Administration, SNMP.
2. Select Configure User from the Step list.
3. Select Add Community from the Action list.
4. Add new community strings as required, and select the corresponding access rights from the Access Mode list.
5. Click Apply

Figure 253: Setting Community Access Strings



To show the community access strings:

1. Click Administration, SNMP.
2. Select Configure User from the Step list.

3. Select Show Community from the Action list.

Figure 254: Showing Community Access Strings

Administration > SNMP		
Step: 5. Configure User	Action: Show Community	
SNMP Community String List Total: 3		
<input type="checkbox"/>	Community String	Access Mode
<input type="checkbox"/>	public	Read-Only
<input type="checkbox"/>	private	Read/Write
<input type="checkbox"/>	spiderman	Read/Write

Delete Revert

Configuring Local SNMPv3 Users

Use the Administration > SNMP (Configure User - Add SNMPv3 Local User) page to authorize management access for SNMPv3 clients, or to identify the source of SNMPv3 trap messages sent from the local switch. Each SNMPv3 user is defined by a unique name. Users must be configured with a specific security level and assigned to a group. The SNMPv3 group restricts users to a specific read, write, and notify view.

Parameters

These parameters are displayed:

- ◆ **User Name** – The name of user connecting to the SNMP agent.
(Range: 1-32 characters)
- ◆ **Group Name** – The name of the SNMP group to which the user is assigned.
(Range: 1-32 characters)
- ◆ **Security Model** – The user security model; SNMP v1, v2c or v3.
- ◆ **Security Level** – The following security levels are only used for the groups assigned to the SNMP security model:
 - **noAuthNoPriv** – There is no authentication or encryption used in SNMP communications. (This is the default security level.)
 - **AuthNoPriv** – SNMP communications use authentication, but the data is not encrypted.
 - **AuthPriv** – SNMP communications use both authentication and encryption.
- ◆ **Authentication Protocol** – The method used for user authentication.
(Options: MD5, SHA; Default: MD5)
- ◆ **Authentication Password** – A minimum of eight plain text characters is required.

- ◆ **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 255: Configuring Local SNMPv3 Users

The screenshot shows a web interface for configuring a local SNMPv3 user. The breadcrumb navigation is "Administration > SNMP". The "Step" dropdown is set to "5. Configure User" and the "Action" dropdown is set to "Add SNMPv3 Local User". The form is divided into several sections:

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

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

To show local SNMPv3 users:

1. Click Administration, SNMP.
2. Select Configure User from the Step list.
3. Select Show SNMPv3 Local User from the Action list.

Figure 256: Showing Local SNMPv3 Users

The screenshot shows the 'Administration > SNMP' configuration page. At the top, there are two dropdown menus: 'Step: 5. Configure User' and 'Action: Show SNMPv3 Local User'. Below this is a table titled 'SNMPv3 Local User List' with a 'Total: 1' indicator. The table has seven columns: 'User Name', 'Group Name', 'Model', 'Level', 'Authentication', and 'Privacy'. A single row is displayed with the following values: 'chris', 'r&d', 'v3', 'authPriv', 'MD5', and 'DES56'. Below the table are two buttons: 'Delete' and 'Revert'.

	User Name	Group Name	Model	Level	Authentication	Privacy
<input type="checkbox"/>	chris	r&d	v3	authPriv	MD5	DES56

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 403](#) and [“Specifying a Remote Engine ID” on page 389](#).)

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 Internet address of the remote device where the user resides.
- ◆ **Security Model** – The user security model; SNMP v1, v2c or v3. (Default: v3)
- ◆ **Security Level** – The following security levels are only used for the groups assigned to the SNMP security model:
 - **noAuthNoPriv** – There is no authentication or encryption used in SNMP communications. (This is the default security level.)
 - **AuthNoPriv** – SNMP communications use authentication, but the data is not encrypted.
 - **AuthPriv** – SNMP communications use both authentication and encryption.

- ◆ **Authentication Protocol** – The method used for user authentication. (Options: MD5, SHA; Default: MD5)
- ◆ **Authentication Password** – A minimum of eight plain text characters is required.
- ◆ **Privacy Protocol** – The encryption algorithm use for data privacy; only 56-bit DES is currently available.
- ◆ **Privacy Password** – A minimum of eight plain text characters is required.

Web Interface

To configure a remote SNMPv3 user:

1. Click Administration, SNMP.
2. Select Configure User from the Step list.
3. Select Add SNMPv3 Remote User from the Action list.
4. Enter a name and assign it to a group. Enter the IP address to identify the source of SNMPv3 inform messages sent from the local switch. If the security model is set to SNMPv3 and the security level is authNoPriv or authPriv, then an authentication protocol and password must be specified. If the security level is authPriv, a privacy password must also be specified.
5. Click Apply

Figure 257: Configuring Remote SNMPv3 Users

The screenshot shows the 'Administration > SNMP' configuration page. At the top, the 'Step' is set to '5. Configure User' and the 'Action' is 'Add SNMPv3 Remote User'. The configuration is organized into sections: 'SNMPv3 User', 'User Authentication', and 'Data Privacy'. Under 'SNMPv3 User', the 'User Name' is 'mark', 'Group Name' has radio buttons for 'public' (selected) and 'r&d', and 'Remote IP' is '192.168.1.19'. Under 'User Authentication', 'Security Model' is 'v3', 'Security Level' is 'authPriv', 'Authentication Protocol' is 'MD5', and 'Authentication Password' is 'greenpeace'. Under 'Data Privacy', 'Privacy Protocol' is 'DESS6' and 'Privacy Password' is 'einstien'. 'Apply' and 'Revert' buttons are at the bottom right.

Field	Value
User Name	mark
Group Name	<input checked="" type="radio"/> public <input type="radio"/> r&d
Remote IP	192.168.1.19
Security Model	v3
Security Level	authPriv
Authentication Protocol	MD5
Authentication Password	greenpeace
Privacy Protocol	DESS6
Privacy Password	einstien

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 258: Showing Remote SNMPv3 Users

The screenshot shows the 'Administration > SNMP' page. At the top, there are two dropdown menus: 'Step: 5. Configure User' and 'Action: Show SNMPv3 Remote User'. Below this is a section titled 'SNMPv3 Remote User List' with a 'Total: 1' indicator. A table contains one row with the following data:

<input type="checkbox"/>	User Name	Group Name	Engine ID	Model	Level	Authentication	Privacy
<input type="checkbox"/>	mark	r&d	5432100000	v3	authPriv	MD5	DES56

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

Specifying Trap Managers

Use the Administration > SNMP (Configure Trap) page to specify the host devices to be sent notifications and the types of notifications to send. Notifications indicating status changes are issued by the switch to the specified notification managers. You must specify notification 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 notification 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 387](#)).
2. Create a view with the required notification messages ([page 390](#)).
3. Configure the group (matching the community string specified on the Configure Trap - Add page) to include the required notify view ([page 393](#)).
4. Enable informs as described in the following pages.

To send an inform to a SNMPv3 host, complete these steps:

1. Enable the SNMP agent ([page 387](#)).

2. Create a local SNMPv3 user to use in the message exchange process (page 399). If the user specified in the notification configuration page does not exist, an SNMPv3 group will be automatically created using the name of the specified local user, and default settings for the read, write, and notify view.
3. Create a view with the required notification messages (page 390).
4. Create a group that includes the required notify view (page 393).
5. Enable informs as described in the following pages.

Parameters

These parameters are displayed:

SNMP Version 1

- ◆ **IP Address** – IP address of a new management station to receive notification message (i.e., the targeted recipient).
- ◆ **Version** – Specifies whether to send notifications using SNMP v1, v2c, or v3. (Default: v1)
- ◆ **Community String** – Specifies a valid community string for the new notification manager entry. (Range: 1-32 characters, case sensitive)
Although you can set this string in the Configure Notification – Add page, we recommend defining it in the Configure User – Add Community page.
- ◆ **UDP Port** – Specifies the UDP port number used by the notification manager. (Default: 162)

SNMP Version 2c

- ◆ **IP Address** – IP address of a new management station to receive notification message (i.e., the targeted recipient).
- ◆ **Version** – Specifies whether to send notifications using SNMP v1, v2c, or v3.
- ◆ **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 notification manager entry. (Range: 1-32 characters, case sensitive)
Although you can set this string in the Configure Notification – Add page, we recommend defining it in the Configure User – Add Community page.
- ◆ **UDP Port** – Specifies the UDP port number used by the notification manager. (Default: 162)

SNMP Version 3

- ◆ **IP Address** – IP address of a new management station to receive notification message (i.e., the targeted recipient).
- ◆ **Version** – Specifies whether to send notifications using SNMP v1, v2c, or v3.
- ◆ **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 notification messages sent from the local switch. (Range: 1-32 characters)
If an account for the specified user has not been created ([page 399](#)), 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 401](#)), one will be automatically generated.
- ◆ **UDP Port** – Specifies the UDP port number used by the notification manager. (Default: 162)

- ◆ **Security Level** – When notification 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 notification 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 259: Configuring Trap Managers (SNMPv1)

The screenshot shows the 'Administration > SNMP' configuration page. At the top, 'Step: 6. Configure Trap' and 'Action: Add' are selected. The form contains the following fields:

- IP Address: 192.168.0.3
- Version: v1
- Community String: private
- UDP Port (1-65535): 162

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

Figure 260: Configuring Trap Managers (SNMPv2c)

The screenshot shows the 'Administration > SNMP' configuration page. At the top, 'Step: 6. Configure Trap' and 'Action: Add' are selected. The form contains the following fields:

- IP Address: 192.168.2.9
- Version: v2c
- Notification Type: Inform
- Timeout (0-2147483647): [empty] centiseconds
- Retry Times (0-255): [empty]
- Community String: venus
- UDP Port (1-65535): [empty]

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

Figure 261: Configuring Trap Managers (SNMPv3)

To show configured notification managers:

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

Figure 262: Showing Notification Managers

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

Creating SNMP Notification Logs

Use the Administration > SNMP (Configure Notify Filter - Add) page to create an SNMP notification log.

Command Usage

- ◆ Systems that support SNMP often need a mechanism for recording Notification information as a hedge against lost notifications, whether there are Traps or Informs that may be exceeding retransmission limits. The Notification Log MIB (NLM, RFC 3014) provides an infrastructure in which information from other MIBs may be logged.

- ◆ Given the service provided by the NLM, individual MIBs can now bear less responsibility to record transient information associated with an event against the possibility that the Notification message is lost, and applications can poll the log to verify that they have not missed any important Notifications.
- ◆ If notification logging is not configured, when the switch reboots, some SNMP traps (such as warm start) cannot be logged.
- ◆ To avoid this problem, notification logging should be configured as described in this section, and these commands stored in the startup configuration file using the System > File (Copy – Running-Config) page as described on [page 73](#). 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 403](#), a default notify filter will be created.

Parameters

These parameters are displayed:

- ◆ **IP Address** – The IPv4 or IPv6 address of a remote device. The specified target host must already have been configured using the Administration > SNMP (Configure Trap – Add) page.

The notification log is stored locally. It is not sent to a remote device. This remote host parameter is only required to complete mandatory fields in the SNMP Notification MIB.

- ◆ **Filter Profile Name** – Notification log profile name. (Range: 1-32 characters)

Web Interface

To create an SNMP notification log:

1. Click Administration, SNMP.
2. Select Configure Notify Filter from the Step list.
3. Select Add from the Action list.
4. Fill in the IP address of a configured trap manager and the filter profile name.
5. Click Apply

Figure 263: Creating SNMP Notification Logs

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 264: Showing SNMP Notification Logs

SNMP Notify Filter List Total: 1		
<input type="checkbox"/>	Filter profile name	IP Address
<input type="checkbox"/>	R&D	192.168.0.99

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 265: Showing SNMP Statistics

Administration > SNMP			
Step: 8. Show Statistics			
SNMP Statistics			
SNMP packets Input	0	SNMP packets Output	0
Bad SNMP version errors	0	Too big errors	0
Unknown community name	0	No such name errors	0
Illegal operation for community name supplied	0	Bad values errors	0
Encoding errors	0	General errors	0
Number of requested variables	0	Response PDUs	0
Number of altered variables	0	Trap PDUs	0
Get-request PDUs	0		
Get-next PDUs	0		
Set-request PDUs	0		

Remote Monitoring

Remote Monitoring allows a remote device to collect information or respond to specified events on an independent basis. This switch is an RMON-capable device which can independently perform a wide range of tasks, significantly reducing network management traffic. It can continuously run diagnostics and log information on network performance. If an event is triggered, it can automatically notify the network administrator of a failure and provide historical information about the event. If it cannot connect to the management agent, it will continue to perform any specified tasks and pass data back to the management station the next time it is contacted.

The switch supports mini-RMON, which consists of the Statistics, History, Event and Alarm groups. When RMON is enabled, the system gradually builds up information about its physical interfaces, storing this information in the relevant RMON database group. A management agent then periodically communicates with the switch using the SNMP protocol. However, if the switch encounters a critical event, it can automatically send a trap message to the management agent which can then respond to the event if so configured.

Configuring RMON Alarms

Use the Administration > RMON (Configure Global - Add - Alarm) page to define specific criteria that will generate response events. Alarms can be set to test data over any specified time interval, and can monitor absolute or changing values (such as a statistical counter reaching a specific value, or a statistic changing by a certain amount over the set interval). Alarms can be set to respond to rising or falling thresholds. (However, note that after an alarm is triggered it will not be triggered again until the statistical value crosses the opposite bounding threshold and then back across the trigger threshold.

Command Usage

- ◆ If an alarm is already defined for an index, the entry must be deleted before any changes can be made.

Parameters

These parameters are displayed:

- ◆ **Index** – Index to this entry. (Range: 1-65535)
- ◆ **Variable** – The object identifier of the MIB variable to be sampled. Only variables of the type etherStatsEntry.n.n may be sampled.

Note that etherStatsEntry.n uniquely defines the MIB variable, and etherStatsEntry.n.n defines the MIB variable, plus the etherStatsIndex. For example, 1.3.6.1.2.1.16.1.1.6.1 denotes etherStatsBroadcastPkts, plus the etherStatsIndex of 1.
- ◆ **Interval** – The polling interval. (Range: 1-31622400 seconds)
- ◆ **Sample Type** – Tests for absolute or relative changes in the specified variable.
 - **Absolute** – The variable is compared directly to the thresholds at the end of the sampling period.
 - **Delta** – The last sample is subtracted from the current value and the difference is then compared to the thresholds.
- ◆ **Rising Threshold** – If the current value is greater than or equal to the rising threshold, and the last sample value was less than this threshold, then an alarm will be generated. After a rising event has been generated, another such event will not be generated until the sampled value has fallen below the rising threshold, reaches the falling threshold, and again moves back up to the rising threshold. (Range: 0-2147483647)
- ◆ **Rising Event Index** – The index of the event to use if an alarm is triggered by monitored variables reaching or crossing above the rising threshold. If there is no corresponding entry in the event control table, then no event will be generated. (Range: 0-65535)
- ◆ **Falling Threshold** – If the current value is less than or equal to the falling threshold, and the last sample value was greater than this threshold, then an alarm will be generated. After a falling event has been generated, another such

event will not be generated until the sampled value has risen above the falling threshold, reaches the rising threshold, and again moves back down to the failing threshold. (Range: 0-2147483647)

- ◆ **Falling Event Index** – The index of the event to use if an alarm is triggered by monitored variables reaching or crossing below the falling threshold. If there is no corresponding entry in the event control table, then no event will be generated. (Range: 0-65535)
- ◆ **Owner** – Name of the person who created this entry. (Range: 1-127 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 266: Configuring an RMON Alarm

The screenshot shows the 'Administration > RMON' configuration page. At the top, there are dropdown menus for 'Step: 1. Configure Global' and 'Action: Add'. Below this, there are radio buttons for 'Alarm' (selected) and 'Event'. The configuration fields are as follows:

Index (1-65535)	1
Variable	6.1
Interval (1-31622400)	15 sec
Sample Type	Delta
Rising Threshold (0-2147483647)	100
Rising Event Index (0-65535)	30
Falling Threshold (0-2147483647)	1
Falling Event Index (0-65535)	2
Owner	bill

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

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 267: Showing Configured RMON Alarms

<input type="checkbox"/>	Index	Status	Variable	Interval	Type	Last Value	Rising Threshold	Rising Event Index	Falling Threshold	Falling Event Index	Owner
<input type="checkbox"/>	1	Valid	1.3.6.1.2.1.16.1.1.1.6.1	30	Delta	0	892800	0	446400	0	
<input type="checkbox"/>	2	Valid	1.3.6.1.2.1.16.1.1.1.6.2	30	Delta	0	892800	0	446400	0	
<input type="checkbox"/>	3	Valid	1.3.6.1.2.1.16.1.1.1.6.3	30	Delta	0	892800	0	446400	0	
<input type="checkbox"/>	4	Valid	1.3.6.1.2.1.16.1.1.1.6.4	30	Delta	0	892800	0	446400	0	
<input type="checkbox"/>	5	Valid	1.3.6.1.2.1.16.1.1.1.6.5	30	Delta	0	892800	0	446400	0	

Configuring RMON Events

Use the Administration > RMON (Configure Global - Add - Event) page to set the action to take when an alarm is triggered. The response can include logging the alarm or sending a message to a trap manager. Alarms and corresponding events provide a way of immediately responding to critical network problems.

Command Usage

- ◆ If an alarm is already defined for an index, the entry must be deleted before any changes can be made.
- ◆ One default event is configured as follows:
 - event Index = 1
 - Description: RMON_TRAP_LOG
 - Event type: log & trap
 - Event community name is public
 - Owner is RMON_SNMP

Parameters

These parameters are displayed:

- ◆ **Index** – Index to this entry. (Range: 1-65535)
- ◆ **Type** – Specifies the type of event to initiate:
 - **None** – No event is generated.

- **Log** – Generates an RMON log entry when the event is triggered. Log messages are processed based on the current configuration settings for event logging (see [“System Log Configuration” on page 355](#)).
 - **Trap** – Sends a trap message to all configured trap managers (see [“Specifying Trap Managers” on page 403](#)).
 - **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 398](#)) prior to configuring it here. (Range: 1-127 characters)
- ◆ **Description** – A comment that describes this event. (Range: 1-127 characters)
- ◆ **Owner** – Name of the person who created this entry. (Range: 1-127 characters)

Web Interface

To configure an RMON event:

1. Click Administration, RMON.
2. Select Configure Global from the Step list.
3. Select Add from the Action list.
4. Click Event.
5. Enter an index number, the type of event to initiate, the community string to send with trap messages, the name of the person who created this event, and a brief description of the event.
6. Click Apply

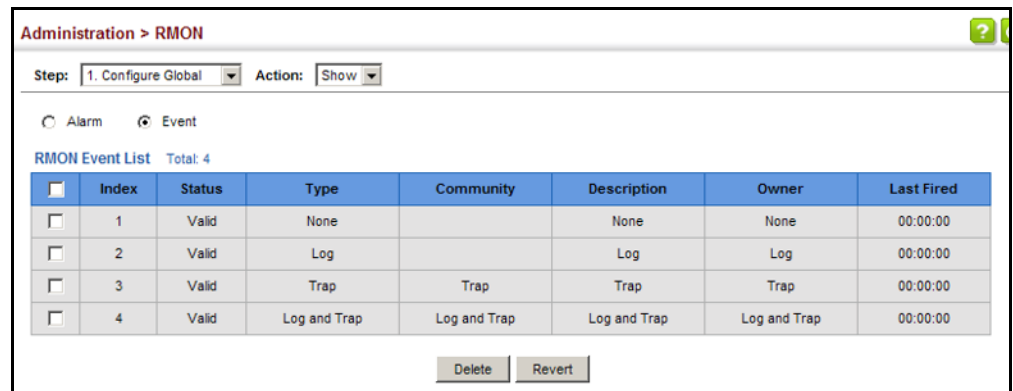
Figure 268: Configuring an RMON Event

The screenshot shows a web interface for configuring an RMON event. The breadcrumb is "Administration > RMON". At the top, there are two dropdown menus: "Step: 1. Configure Global" and "Action: Add". Below these are two radio buttons: "Alarm" (unselected) and "Event" (selected). The main configuration area contains several labeled input fields: "Index (1-65535)" with the value "2", "Type" with a dropdown menu set to "Log and Trap", "Community" with the value "private", "Description" with the value "for software group", and "Owner" with the value "david". At the bottom right of the form are two buttons: "Apply" and "Revert".

To show configured RMON events:

1. Click Administration, RMON.
2. Select Configure Global from the Step list.
3. Select Show from the Action list.
4. Click Event.

Figure 269: Showing Configured RMON Events



The screenshot shows the 'Administration > RMON' configuration page. At the top, there is a breadcrumb 'Administration > RMON' and a help icon. Below that, there are two dropdown menus: 'Step: 1. Configure Global' and 'Action: Show'. There are two radio buttons: 'Alarm' (unselected) and 'Event' (selected). Below the radio buttons is the heading 'RMON Event List Total: 4'. A table with 8 columns is displayed: Index, Status, Type, Community, Description, Owner, and Last Fired. The table contains four rows of data. At the bottom of the table are two buttons: 'Delete' and 'Revert'.

<input type="checkbox"/>	Index	Status	Type	Community	Description	Owner	Last Fired
<input type="checkbox"/>	1	Valid	None		None	None	00:00:00
<input type="checkbox"/>	2	Valid	Log		Log	Log	00:00:00
<input type="checkbox"/>	3	Valid	Trap	Trap	Trap	Trap	00:00:00
<input type="checkbox"/>	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 105](#).
- ◆ The switch reserves two index entries for each port. If a default index entry is re-assigned to another port using the Add page, this index will not appear in the Show nor Show Details page for the port to which is normally assigned. For

example, if control entry 15 is assigned to port 5, this index entry will be removed from the Show and Show Details page for port 8.

Parameters

These parameters are displayed:

- ◆ **Port** – The port number on the switch.
- ◆ **Index** - Index to this entry. (Range: 1-65535)
- ◆ **Interval** - The polling interval. (Range: 1-3600 seconds; Default: 1800 seconds)
- ◆ **Buckets** - The number of buckets requested for this entry. (Range: 1-65536; Default: 50)
The number of buckets granted are displayed on the Show page.
- ◆ **Owner** - Name of the person who created this entry. (Range: 1-127 characters)

Web Interface

To periodically sample statistics on a port:

1. Click Administration, RMON.
2. Select Configure Interface from the Step list.
3. Select Add from the Action list.
4. Click History.
5. Select a port from the list as the data source.
6. Enter an index number, the sampling interval, the number of buckets to use, and the name of the owner for this entry.
7. Click Apply

Figure 270: Configuring an RMON History Sample

The screenshot shows the 'Administration > RMON' configuration page. At the top, there are two dropdown menus: 'Step: 2. Configure Interface' and 'Action: Add'. Below these are two radio buttons: 'History' (selected) and 'Statistics'. The 'Port' is set to '2'. The 'Index (1-65535)' is '100', 'Interval (1-3600)' is '60' with 'sec' next to it, 'Buckets (1-65535)' is '10', and 'Owner' is 'david'. At the bottom right, there are 'Apply' and 'Revert' buttons.

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 271: Showing Configured RMON History Samples

Administration > RMON

Step: 2. Configure Interface Action: Show

History Statistics

Port 1

RMON History Port List Total: 2

<input type="checkbox"/>	Index	Status	Interval	Requested Buckets	Granted Buckets	Owner
<input type="checkbox"/>	1	Valid	1800	8	8	
<input type="checkbox"/>	2	Valid	30	8	8	

Delete Revert

To show collected RMON history samples:

1. Click Administration, RMON.
2. Select Configure Interface from the Step list.
3. Select Show Details from the Action list.
4. Select a port from the list.
5. Click History.

Figure 272: Showing Collected RMON History Samples

Administration > RMON

Step: 2. Configure Interface Action: Show Details

History Statistics

Port 1

RMON History Details Port List Total: 9

History Index	Sample Index	Interval Start	Octets	Packets	Broadcast Packets	Multicast Packets	Undersize Packets	Oversize Packets	Fragments	Jabbers	CRC Align Errors	Collisions	Drop Events	Network Utilization
1	1	00:00:01	756105	3218	91	894	0	0	0	0	0	0	0	0
2	71	00:35:01	21490	76	0	15	0	0	0	0	0	0	0	0
2	72	00:35:31	46521	120	0	15	0	0	0	0	0	0	0	0
2	73	00:36:01	21682	79	0	15	0	0	0	0	0	0	0	0
2	74	00:36:31	21554	77	0	15	0	0	0	0	0	0	0	0

Configuring RMON Statistical Samples Use the Administration > RMON (Configure Interface - Add - Statistics) page to collect statistics on a port, which can subsequently be used to monitor the network for common errors and overall traffic rates.

Command Usage

- ◆ If statistics collection is already enabled on an interface, the entry must be deleted before any changes can be made.
- ◆ The information collected for each entry includes:
input octets, packets, broadcast packets, multicast packets, undersize packets, oversize packets, CRC alignment errors, jabbers, fragments, collisions, drop events, and frames of various sizes.

Parameters

These parameters are displayed:

- ◆ **Port** – The port number on the switch.
- ◆ **Index** - Index to this entry. (Range: 1-65535)
- ◆ **Owner** - Name of the person who created this entry. (Range: 1-127 characters)

Web Interface

To enable regular sampling of statistics on a port:

1. Click Administration, RMON.
2. Select Configure Interface from the Step list.
3. Select Add from the Action list.
4. Click Statistics.
5. Select a port from the list as the data source.
6. Enter an index number, and the name of the owner for this entry
7. Click Apply

Figure 273: Configuring an RMON Statistical Sample

The screenshot shows the 'Administration > RMON' configuration page. At the top, the breadcrumb 'Administration > RMON' is visible. Below it, the 'Step' is set to '2. Configure Interface' and the 'Action' is set to 'Add'. There are two radio buttons: 'History' (unselected) and 'Statistics' (selected). The 'Port' is set to '2'. The 'Index (1-65535)' is set to '100' and the 'Owner' is set to '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 274: Showing Configured RMON Statistical Samples

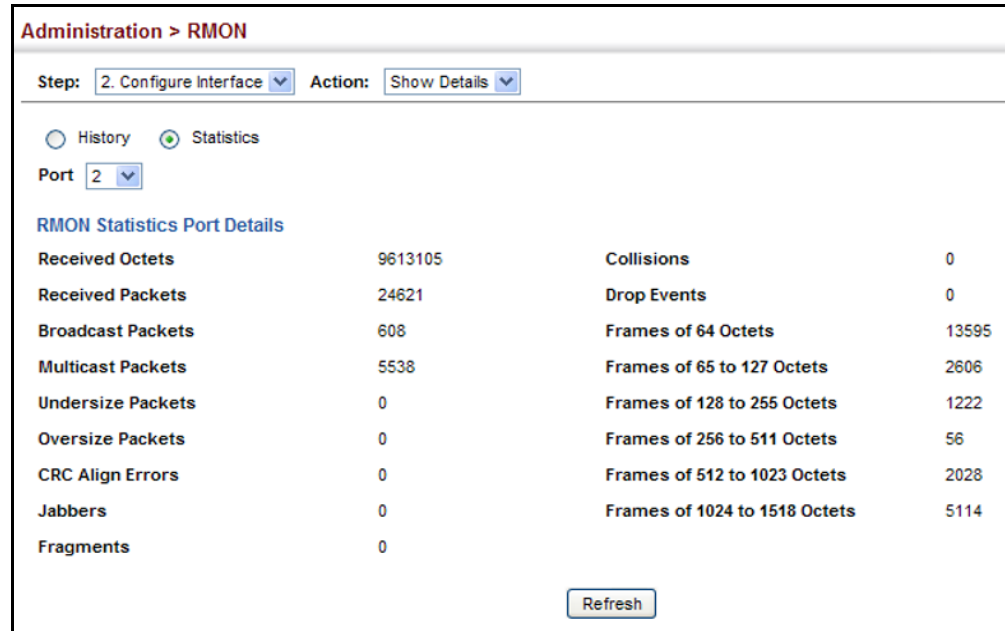
The screenshot shows the 'Administration > RMON' configuration page. At the top, the breadcrumb 'Administration > RMON' is visible. Below it, the 'Step' is set to '2. Configure Interface' and the 'Action' is set to 'Show'. There are two radio buttons: 'History' (unselected) and 'Statistics' (selected). The 'Port' is set to '2'. Below the form, there is a table titled 'RMON Statistics Port List' with a total of 2 samples. The table has columns for 'Index', 'Status', and 'Owner'. At the bottom right, there are 'Delete' and 'Revert' buttons.

	Index	Status	Owner
<input type="checkbox"/>	1	Valid	abc
<input type="checkbox"/>	2	Valid	test

To show collected RMON statistical samples:

1. Click Administration, RMON.
2. Select Configure Interface from the Step list.
3. Select Show Details from the Action list.
4. Select a port from the list.
5. Click Statistics.

Figure 275: 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.

- ◆ 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 276: Configuring a Switch Cluster

Administration > Cluster

Step: 1. Configure Global

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

Apply Revert

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

Parameters

These parameters are displayed:

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

Web Interface

To configure cluster members:

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

Figure 277: Configuring Cluster Members

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 278: Showing Cluster Members

Member ID	Role	IP Address	MAC Address	Description
1	Active Member	10.254.254.2	11-22-33-44-55-33	Managed 24GPoE+2GT+2GSFP Switch
2	Candidate	10.254.254.3	11-22-33-44-55-77	Managed 24GPoE+2GT+2GSFP Switch

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 279: Showing Cluster Candidates

Role	MAC Address	Description
Candidate	11-22-33-44-55-11	Managed 24GPoE+2GT+2GSFP Switch
Active Member	11-22-33-44-55-22	Managed 24GPoE+2GT+2GSFP Switch
Candidate	11-22-33-44-55-33	Managed 24GPoE+2GT+2GSFP Switch
Candidate	11-22-33-44-55-44	Managed 24GPoE+2GT+2GSFP Switch

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.

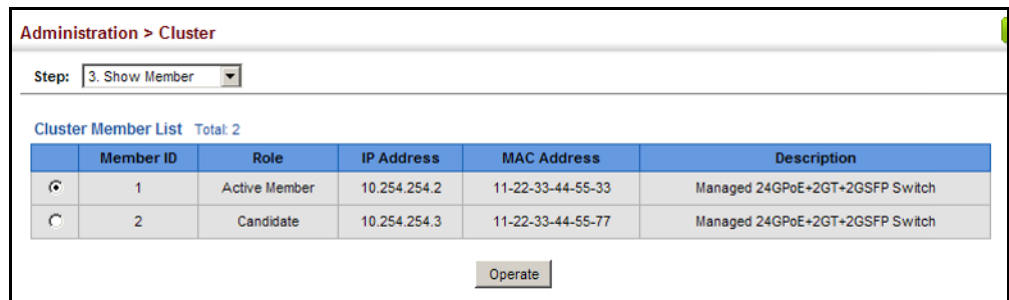
Operate – Remotely manage a cluster member.

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 280: Managing a Cluster Member



IP Configuration

This chapter describes how to configure an IP interface for management access to the switch over the network. This switch supports both IP Version 4 and Version 6, and can be managed simultaneously through either of these address types. You can manually configure a specific IPv4 or IPv6 address, or direct the switch to obtain an IPv4 address using Auto IP, or from a BOOTP or DHCP server when it is powered on. An IPv6 address can either be manually configured or dynamically generated.

This chapter provides information on network functions including:

- ◆ [Ping](#) – Sends ping message to another node on the network.
- ◆ [Address Resolution Protocol](#) – Specifies the timeout for ARP cache entries. Also shows how to display the ARP cache.
- ◆ [IPv4 Configuration](#) – Sets an IPv4 address for management access.
- ◆ [IPv6 Configuration](#) – Sets an IPv6 address for management access.

Using the Ping Function

Use the IP > General > Ping page to send ICMP echo request packets to another node on the network.

Parameters

These parameters are displayed:

- ◆ **Host Name/IP Address** – Specifies the host name (that is, alias) or IPv4/IPv6 address of the target.

For host name-to-IP address translation to function properly, host name lookup must be enabled ([“Configuring General DNS Service Parameters” on page 453](#)), and one or more DNS servers specified (see [“Configuring a List of Name Servers” on page 456](#), or [“Configuring Static DNS Host to Address Entries” on page 457](#)).

- ◆ **Probe Count** – Number of packets to send. (Range: 1-16)
- ◆ **Packet Size** – Number of bytes in a packet. (Range: 0-1452 bytes)

The actual packet size will be eight bytes larger than the size specified because the switch adds header information.

Command Usage

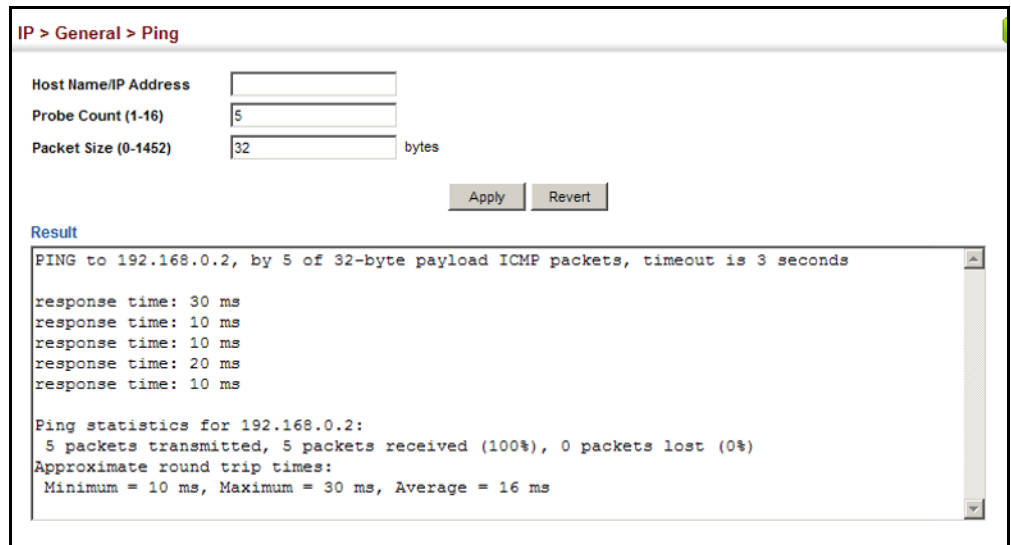
- ◆ Use the ping command to see if another site on the network can be reached.
- ◆ The following are some results of the **ping** command:
 - *Normal response* - The normal response occurs in one to ten seconds, depending on network traffic.
 - *Destination does not respond* - If the host does not respond, a “timeout” appears in ten seconds.
 - *Destination unreachable* - The gateway for this destination indicates that the destination is unreachable.
 - *Network or host unreachable* - The gateway found no corresponding entry in the route table.

Web Interface

To ping another device on the network:

1. Click IP, General, Ping.
2. Specify the target device and ping parameters.
3. Click Apply.

Figure 281: Pinging a Network Device



Address Resolution Protocol

The switch 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 switch (or any standards-based switch/router), it first looks up the MAC address corresponding to the destination IP address in the ARP cache. If the address is found, the switch 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 switch 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 31: 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 switch 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.

Setting the ARP Timeout

Use the IP > ARP (Configure General) page to specify the timeout for ARP cache entries.

Parameters

These parameters are displayed:

- ◆ **Timeout** – Sets the aging time for dynamic entries in the ARP cache. (Range: 300 - 86400 seconds; Default: 1200 seconds or 20 minutes)

The ARP aging timeout can only be set globally for all VLANs.

The aging time determines how long dynamic entries remain in the cache. If the timeout is too short, the switch may tie up resources by repeating ARP requests for addresses recently flushed from the table.

When a ARP entry expires, it is deleted from the cache and an ARP request packet is sent to re-establish the MAC address.

Web Interface

To configure the timeout for the ARP cache:

1. Click IP, ARP.
2. Select Configure General from the Step List.
3. Set the timeout to a suitable value for the ARP cache.
4. Click Apply.

Figure 282: Setting the ARP Timeout

The screenshot shows the 'IP > ARP' configuration page. The 'Step' dropdown is set to '1. Configure General'. The 'Timeout (300-86400)' field is set to '1200' with the unit 'sec'. There are 'Apply' and 'Revert' buttons at the bottom right.

Displaying ARP Entries Use the IP > ARP (Show Information) page to display dynamic or local entries in the ARP cache. The ARP cache contains 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 entries in the ARP cache:

1. Click IP, ARP.
2. Select Show Information from the Step List.

Figure 283: Displaying ARP Entries

The screenshot shows the 'IP > ARP' configuration page with the 'Step' dropdown set to '2. Show Information'. Below the configuration fields, there is a table titled 'Dynamic Address List Total: 1'. The table has three columns: 'IP Address', 'MAC Address', and 'Interface'. The entry shown is IP Address: 192.168.0.99, MAC Address: 00-E0-29-94-34-64, and Interface: VLAN 1. There is a 'Clear' button at the bottom.

IP Address	MAC Address	Interface
192.168.0.99	00-E0-29-94-34-64	VLAN 1

Setting the Switch's IP Address (IP Version 4)

Use the System > IP page to configure an IPv4 address for management access over the network. 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 434](#).

You can direct the device to obtain an address through Auto IP, 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.

To configure an address compatible with your network, you may need to change the switch's default settings. You may also need to establish a default gateway between the switch and management stations that exist on another network segment.

Parameters

These parameters are displayed:

- ◆ **Management 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 (Static), Auto IP, Dynamic Host Configuration Protocol (DHCP), or Boot Protocol (BOOTP). (Default: DHCP)
 - 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.
 - If IP address mode is set to Auto IP, the switch randomly selects an IPv4 link-local address from 169.254.0.1~169.254.255.254. Before starting to use it, the switch tests to see if the address is already in use.

Conflict detection is done using ARP probes. The switch waits a predefined number of seconds after the transmission of the last ARP probe. If no conflicting ARP Reply or ARP Probe has been received, then the switch has successfully claimed the desired IPv4 link-local address.

When a network interface transitions from inactive to active state, the switch does not know what IPv4 link-local addresses are currently in use on that link since the point of attachment may have changed or the network interface may have been inactive when a conflicting address was claimed. In this case, the switch again starts sending ARP probes to see if the previously selecting address has been claimed by another device on the network.

- ◆ **IP Address** – Address of the VLAN to which the management station is attached. Valid IP addresses consist of four numbers, 0 to 255, separated by periods. (Default: None)
- ◆ **Subnet Mask** – This mask identifies the host address bits used for routing to specific subnets. (Default: None)
- ◆ **Gateway IP Address** – IP address of the gateway router between the switch and management stations that exist on other network segments. (Default: 0.0.0.0)
- ◆ **MAC Address** – The physical layer address for this switch.
- ◆ **Restart DHCP** – Requests a new IP address from the DHCP server.

Web Interface

To set a static address for the switch:

1. Click System, IP.
2. Select the VLAN through which the management station is attached, set the IP Address Mode to “Static,” enter the IP address, subnet mask and gateway.
3. Click Apply.

Figure 284: Configuring a Static IPv4 Address

The screenshot shows the 'System > IP' configuration page. It features several fields for configuration: 'Management VLAN' is set to 1; 'IP Address Mode' is set to 'Static'; 'IP Address' is 192.168.0.93; 'Subnet Mask' is 255.255.255.0; 'Gateway IP Address' is 192.168.0.1; and 'MAC Address' is 00-00-00-00-00-02. At the bottom, there is a 'Restart DHCP' button with a tooltip that says 'Click this button to restart DHCP service.', and 'Apply' and 'Revert' buttons.

To obtain an dynamic address through Auto IP for the switch:

1. Click System, IP.
2. Select the VLAN through which the management station is attached, set the IP Address Mode to “Auto IP.”
3. Click Apply to save your changes.
4. Click Apply.

Figure 285: Configuring an Auto IP Address

System > IP

Management VLAN	1
IP Address Mode	Auto IP
IP Address	192.168.0.3
Subnet Mask	255.255.255.0
Gateway IP Address	
MAC Address	00-E0-0C-00-00-FD

Restart DHCP [Click this button to restart DHCP service.](#)

Apply Revert

To obtain an dynamic address through DHCP/BOOTP for the switch:

1. Click System, IP.
2. Select the VLAN through which the management station is attached, set the IP Address Mode to “DHCP” or “BOOTP.”
3. Click Apply to save your changes.
4. Then click Restart DHCP to immediately request a new address.

Figure 286: Configuring a Dynamic IPv4 Address

System > IP

Management VLAN	1
IP Address Mode	DHCP
IP Address	192.168.0.93
Subnet Mask	255.255.255.0
Gateway IP Address	192.168.0.1
MAC Address	00-00-00-00-00-02

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.

Setting the Switch's IP Address (IP Version 6)

This section describes how to configure an IPv6 interface for management access over the network. 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 431](#).

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).

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:

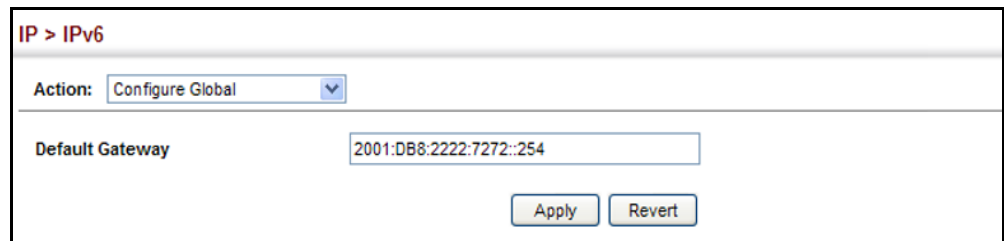
- ◆ **Default Gateway** – Sets the IPv6 address of the default next hop router.
 - An IPv6 default gateway must be defined if the management station is located in a different IPv6 segment.
 - An IPv6 default gateway can only be successfully set when a network interface that directly connects to the gateway has been configured on the switch.

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 287: 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 for "Action" with "Configure Global" selected. Underneath, there is a field for "Default Gateway" containing the IPv6 address "2001:DB8:2222:7272::254". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

Configuring IPv6 Interface Settings

Use the IP > IPv6 Configuration (Configure Interface) page to configure general IPv6 settings for the selected VLAN, including auto-configuration of a global unicast interface address, 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 always 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 will also create a link-local address, but will not generate a global IPv6 address if auto-configuration is not enabled. In this case, you must manually configure an address (see ["Configuring an IPv6 Address" on page 440](#)).
- ◆ 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. By default, all ports on the switch are members of VLAN 1. However, the management station can be attached to a port belonging to any VLAN, as long as that VLAN has been assigned an IP address. (Range: 1-4094)

- ◆ **Address Autoconfig** – Enables stateless autoconfiguration of an IPv6 address on an interface and enables IPv6 functionality on that interface. The network portion of the address is based on prefixes received in IPv6 router advertisement messages, and the host portion is automatically generated using the modified EUI-64 form of the interface identifier (i.e., the switch's MAC address).
 - If a link local address has not yet been assigned to this interface, this command will dynamically generate one. The link-local address is made with an address prefix in the range of FE80~FEBF and a host portion based the switch's MAC address in modified EUI-64 format. It will also generate a global unicast address if a global prefix is included in received router advertisements.
 - When DHCPv6 is 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 will attempt to acquire other non-address configuration information (such as a default gateway).
 - If auto-configuration is not selected, then an address must be manually configured using the Add Interface page described below.

- ◆ **Enable IPv6 Explicitly** – Enables IPv6 on an interface. 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.
 - IPv6 routers do not fragment IPv6 packets forwarded from other routers. However, traffic originating from an end-station connected to an IPv6 router may be fragmented.
 - All devices on the same physical medium must use the same MTU in order to operate correctly.

- IPv6 must be enabled on an interface before the MTU can be set. If an IPv6 address has not been assigned to the switch, "N/A" is displayed in the MTU field.
- ◆ **ND DAD Attempts** – The number of consecutive neighbor solicitation messages sent on an interface during duplicate address detection. (Range: 0-600, Default: 1)
 - Configuring a value of 0 disables duplicate address detection.
 - Duplicate address detection determines if a new unicast IPv6 address already exists on the network before it is assigned to an interface.
 - Duplicate address detection is stopped on any interface that has been suspended (see ["Configuring VLAN Groups" on page 138](#)). 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.
- ◆ **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)
- ◆ **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 Mode

◆ **Interface** – Shows port or trunk configuration page.

◆ **RA Guard** – Blocks incoming Router Advertisement and Router Redirect packets. (Default: Disabled)

IPv6 Router Advertisements (RA) convey information that enables nodes to auto-configure on the network. This information may include the default router address taken from the observed source address of the RA message, as well as on-link prefix information. However, note that unintended misconfigurations, or possibly malicious attacks on the network, may lead to bogus RAs being sent, which in turn can cause operational problems for hosts on the network.

RA Guard can be used to block RAs and Router Redirect (RR) messages on the specified interface. Determine which interfaces are connected to known routers, and enable RA Guard on all other untrusted interfaces.

Web Interface

To general IPv6 settings for the switch:

1. Click IP, IPv6 Configuration.
2. Select Configure Interface from the Action list.
3. Specify the VLAN to configure, enable address auto-configuration, or enable IPv6 explicitly to automatically configure a link-local address and enable IPv6

on the selected interface. Set the MTU size, the maximum number of duplicate address detection messages, the neighbor solicitation message interval, and the remote node reachable time.

4. Click Apply.

Figure 288: Configuring General Settings for an IPv6 Interface

The screenshot shows the 'IP > IPv6 Configuration' page. The 'Action' dropdown is set to 'Configure Interface'. Under 'Mode', 'VLAN' is selected with a radio button. The 'VLAN' dropdown is set to '1'. The 'Address Autoconfig' and 'Enable IPv6 Explicitly' checkboxes are both unchecked. The 'MTU (1280-65535)' is set to '1500' bytes. The 'ND DAD Attempts (0-600)' is set to '1'. The 'ND NS Interval (1000-3600000)' is set to '1000' ms. The 'ND Reachable-Time (0-3600000)' is set to '30000' ms. There is a 'Restart DHCPv6' button with a tooltip that says 'Click this button to restart DHCPv6 service.' At the bottom right are 'Apply' and 'Revert' buttons.

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 289: 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 with a radio button. Under 'Interface', 'Port' is selected with a radio button. Below this is a 'Port List' table with 28 total ports. The table has columns for 'Port' and 'RA Guard'. The first five rows are visible, showing ports 1 through 5. Port 5 has the 'RA Guard' checkbox checked, while ports 1-4 have it unchecked.

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

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 435](#)), 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 the 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 435](#)).
 - 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. 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.

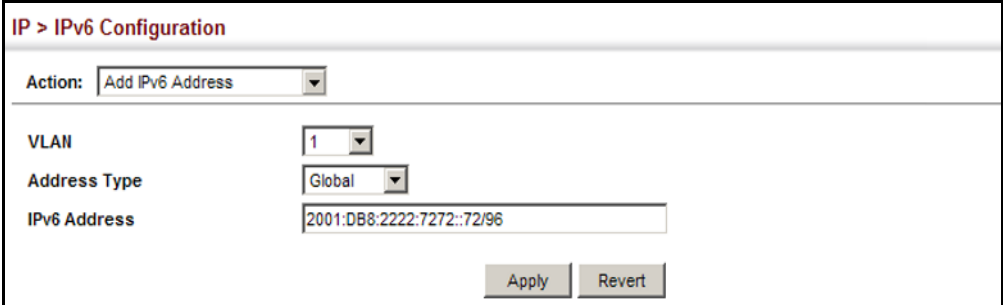
- **Link Local** – Configures an IPv6 link-local address.
 - The address prefix must be in the range of FE80~FEBF.
 - You can configure only one link-local address per interface.
 - The specified address replaces a link-local address that was automatically generated for the interface.
- ◆ **IPv6 Address** – 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 290: Configuring an IPv6 Address



The screenshot shows the 'IP > IPv6 Configuration' web interface. At the top, there is a breadcrumb 'IP > IPv6 Configuration'. Below it, the 'Action' dropdown menu is set to 'Add IPv6 Address'. The 'VLAN' dropdown is set to '1', the 'Address Type' dropdown is set to 'Global', and the 'IPv6 Address' text field contains '2001:DB8:2222:7272::72/96'. At the bottom right, there are 'Apply' and 'Revert' buttons.

Showing IPv6 Addresses Use the IP > IPv6 Configuration (Show IPv6 Address) page to display the IPv6 addresses assigned to an interface.

Parameters

These parameters are displayed:

- ◆ **VLAN** – ID of a configured VLAN which is to be used for management access. 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)
- ◆ **IP Address Type** – The address type (Global, EUI-64, Link Local).
- ◆ **IP Address** – An IPv6 address assigned to this interface.

In addition to the unicast addresses assigned to an interface, a host is also required to listen to the all-nodes multicast addresses FF01::1 (interface-local scope) and FF02::1 (link-local scope).

FF01::1/16 is the transient interface-local multicast address for all attached IPv6 nodes, and FF02::1/16 is the link-local multicast address for all attached IPv6 nodes. The interface-local multicast address is only used for loopback transmission of multicast traffic. Link-local multicast addresses cover the same types as used by link-local unicast addresses, including all nodes (FF02::1), all routers (FF02::2), and solicited nodes (FF02::1:FFXX:XXXX) as described below.

A node is also required to compute and join the associated solicited-node multicast addresses for every unicast and anycast address it is assigned. IPv6 addresses that differ only in the high-order bits, e.g. due to multiple high-order prefixes associated with different aggregations, will map to the same solicited-node address, thereby reducing the number of multicast addresses a node must join. In this example, FF02::1:FF90:0/104 is the solicited-node multicast address which is formed by taking the low-order 24 bits of the address and appending those bits to the prefix.

Note that the solicited-node multicast address (link-local scope FF02) is used to resolve the MAC addresses for neighbor nodes since IPv6 does not support the broadcast method used by the Address Resolution Protocol in IPv4.

These additional addresses are displayed by the CLI (see the “show ip interface” command in the *CLI Reference Guide*).

- ◆ **Configuration Mode** – Indicates if this address was automatically generated for 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 291: Showing Configured IPv6 Addresses

The screenshot shows a web interface for IP > IPv6 configuration. At the top, there is a breadcrumb 'IP > IPv6'. Below it, an 'Action:' dropdown menu is set to 'Show IPv6 Address'. Underneath, a 'VLAN' dropdown menu is set to '1'. The main content area is titled 'IPv6 Address List Total: 3'. It contains a table with the following data:

<input type="checkbox"/>	IP Address Type	IP Address	Configuration Mode
<input type="checkbox"/>	EUI-64	2001:DB8::1:200:E8FF:FE93:82A0/64	Manual
<input type="checkbox"/>	Global	2001:DB8:2222:7272::72/96	Manual
<input type="checkbox"/>	Link Local	FE80::200:E8FF:FE93:82A0%1/64	Auto

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

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 32: Show IPv6 Neighbors - display description

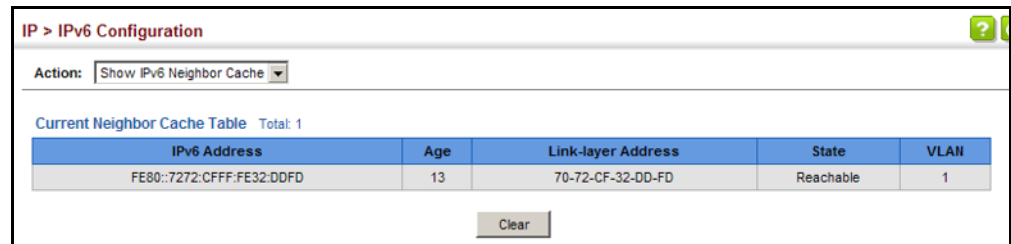
Field	Description
IPv6 Address	IPv6 address of neighbor
Age	The time since the address was verified as reachable (in seconds). A static entry is indicated by the value "Permanent."
Link-layer Addr	Physical layer MAC address.
State	<p>The following states are used for dynamic entries:</p> <ul style="list-style-type: none">◆ Incomplete - Address resolution is being carried out on the entry. A neighbor solicitation message has been sent to the multicast address of the target, but it has not yet returned a neighbor advertisement message.◆ Invalid - An invalidated mapping. Setting the state to invalid dis-associates the interface identified with this entry from the indicated mapping (RFC 4293).◆ Reachable - Positive confirmation was received within the last ReachableTime interval that the forward path to the neighbor was functioning. While in REACH state, the device takes no special action when sending packets.◆ Stale - More than the ReachableTime interval has elapsed since the last positive confirmation was received that the forward path was functioning. While in STALE state, the device takes no action until a packet is sent.◆ 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 resending neighbor solicitation messages every RetransTimer interval until confirmation of reachability is received.◆ Unknown - Unknown state. <p>The following states are used for static entries:</p> <ul style="list-style-type: none">◆ Incomplete -The interface for this entry is down.◆ 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 292: Showing IPv6 Neighbors



The screenshot shows the 'IP > IPv6 Configuration' page. At the top, there is a breadcrumb trail and a help icon. Below that, an 'Action:' dropdown menu is set to 'Show IPv6 Neighbor Cache'. Underneath, it says 'Current Neighbor Cache Table Total: 1'. A table with five columns is displayed: 'IPv6 Address', 'Age', 'Link-layer Address', 'State', and 'VLAN'. The table contains one row with the following data: IPv6 Address: FE80::7272:CFFF:FE32:DDFD, Age: 13, Link-layer Address: 70-72-CF-32-DD-FD, State: Reachable, and VLAN: 1. A 'Clear' button is located at the bottom right of the table area.

IPv6 Address	Age	Link-layer Address	State	VLAN
FE80::7272:CFFF:FE32:DDFD	13	70-72-CF-32-DD-FD	Reachable	1

Showing IPv6 Statistics Use the IP > IPv6 Configuration (Show Statistics) page to display statistics about IPv6 traffic passing through this switch.

Command Usage

This switch provides statistics for the following traffic types:

- ◆ **IPv6** – The Internet Protocol for Version 6 addresses provides a mechanism for transmitting blocks of data (often called packets or frames) from a source to a destination, where these network devices (that is, hosts) are identified by fixed length addresses. The Internet Protocol also provides for fragmentation and reassembly of long packets, if necessary, for transmission through “small packet” networks.
- ◆ **ICMPv6** – Internet Control Message Protocol for Version 6 addresses is a network layer protocol that transmits message packets to report errors in processing IPv6 packets. ICMP is therefore an integral part of the Internet Protocol. ICMP messages may be used to report various situations, such as when a datagram cannot reach its destination, when the gateway does not have the buffering capacity to forward a datagram, and when the gateway can direct the host to send traffic on a shorter route. ICMP is also used by routers to feed back information about more suitable routes (that is, the next hop router) to use for a specific destination.
- ◆ **UDP** – User Datagram Protocol provides a datagram mode of packet switched communications. It uses IP as the underlying transport mechanism, providing access to IP-like services. UDP packets are delivered just like IP packets – connection-less datagrams that may be discarded before reaching their targets. UDP is useful when TCP would be too complex, too slow, or just unnecessary.

Parameters

These parameters are displayed:

Table 33: 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.
Reassembly Succeeded	The number of IPv6 datagrams successfully reassembled. Note that this counter is incremented at the interface to which these datagrams were addressed which might not be necessarily the input interface for some of the fragments.
Reassembly Failed	The number of failures detected by the IPv6 re-assembly algorithm (for whatever reason: timed out, errors, etc.). Note that this is not necessarily a count of discarded IPv6 fragments since some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by combining them as they are received. This counter is incremented at the interface to which these fragments were addressed which might not be necessarily the input interface for some of the fragments.

Table 33: Show IPv6 Statistics - display description (Continued)

Field	Description
<i>IPv6 Transmitted</i>	
Forwards Datagrams	The number of output datagrams which this entity received and forwarded to their final destinations. In entities which do not act as IPv6 routers, this counter will include only those packets which were Source-Routed via this entity, and the Source-Route processing was successful. Note that for a successfully forwarded datagram the counter of the outgoing interface is incremented.
Requests	The total number of IPv6 datagrams which local IPv6 user-protocols (including ICMP) supplied to IPv6 in requests for transmission. Note that this counter does not include any datagrams counted in <code>ipv6IfStatsOutForwDatagrams</code> .
Discards	The number of output IPv6 datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (e.g., for lack of buffer space). Note that this counter would include datagrams counted in <code>ipv6IfStatsOutForwDatagrams</code> if any such packets met this (discretionary) discard criterion.
No Routes	The number of input datagrams discarded because no route could be found to transmit them to their destination.
Generated Fragments	The number of output datagram fragments that have been generated as a result of fragmentation at this output interface.
Fragment Succeeded	The number of IPv6 datagrams that have been successfully fragmented at this output interface.
Fragment Failed	The number of IPv6 datagrams that have been discarded because they needed to be fragmented at this output interface but could not be.
ICMPv6 Statistics	
<i>ICMPv6 received</i>	
Input	The total number of ICMP messages received by the interface which includes all those counted by <code>ipv6IfIcmpInErrors</code> . Note that this interface is the interface to which the ICMP messages were addressed which may not be necessarily the input interface for the messages.
Errors	The number of ICMP messages which the interface received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, etc.).
Destination Unreachable Messages	The number of ICMP Destination Unreachable messages received by the interface.
Packet Too Big Messages	The number of ICMP Packet Too Big messages received by the interface.
Time Exceeded Messages	The number of ICMP Time Exceeded messages received by the interface.
Parameter Problem Messages	The number of ICMP Parameter Problem messages received by the interface.
Echo Request Messages	The number of ICMP Echo (request) messages received by the interface.
Echo Reply Messages	The number of ICMP Echo Reply messages received by the interface.
Router Solicit Messages	The number of ICMP Router Solicit messages received by the interface.
Router Advertisement Messages	The number of ICMP Router Advertisement messages received by the interface.
Neighbor Solicit Messages	The number of ICMP Neighbor Solicit messages received by the interface.

Table 33: Show IPv6 Statistics - display description (Continued)

Field	Description
Neighbor Advertisement Messages	The number of ICMP Neighbor Advertisement messages received by the interface.
Redirect Messages	The number of Redirect messages received by the interface.
Group Membership Query Messages	The number of ICMPv6 Group Membership Query messages received by the interface.
Group Membership Response Messages	The number of ICMPv6 Group Membership Response messages received by the interface.
Group Membership Reduction Messages	The number of ICMPv6 Group Membership Reduction messages received by the interface.
Multicast Listener Discovery Version 2 Reports	The number of MLDv2 reports received by the interface.
<i>ICMPv6 Transmitted</i>	
Output	The total number of ICMP messages which this interface attempted to send. Note that this counter includes all those counted by icmpOutErrors.
Destination Unreachable Messages	The number of ICMP Destination Unreachable messages sent by the interface.
Packet Too Big Messages	The number of ICMP Packet Too Big messages sent by the interface.
Time Exceeded Messages	The number of ICMP Time Exceeded messages sent by the interface.
Parameter Problem Message	The number of ICMP Parameter Problem messages sent by the interface.
Echo Request Messages	The number of ICMP Echo (request) messages sent by the interface.
Echo Reply Messages	The number of ICMP Echo Reply messages sent by the interface.
Router Solicit Messages	The number of ICMP Router Solicitation messages sent by the interface.
Router Advertisement Messages	The number of ICMP Router Advertisement messages sent by the interface.
Neighbor Solicit Messages	The number of ICMP Neighbor Solicit messages sent by the interface.
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 Response Messages	The number of ICMPv6 Group Membership Response messages sent.
Group Membership Reduction Messages	The number of ICMPv6 Group Membership Reduction messages sent.
Multicast Listener Discovery Version 2 Reports	The number of MLDv2 reports sent by the interface.
UDP Statistics	
Input	The total number of UDP datagrams delivered to UDP users.
No Port Errors	The total number of received UDP datagrams for which there was no application at the destination port.

Table 33: Show IPv6 Statistics - display description (Continued)

Field	Description
Other Errors	The number of received UDP datagrams that could not be delivered for reasons other than the lack of an application at the destination port.
Output	The total number of UDP datagrams sent from this entity.

Web Interface

To show the IPv6 statistics:

1. Click IP, IPv6 Configuration.
2. Select Show Statistics from the Action list.
3. Click IPv6, ICMPv6 or UDP.

Figure 293: Showing IPv6 Statistics (IPv6)

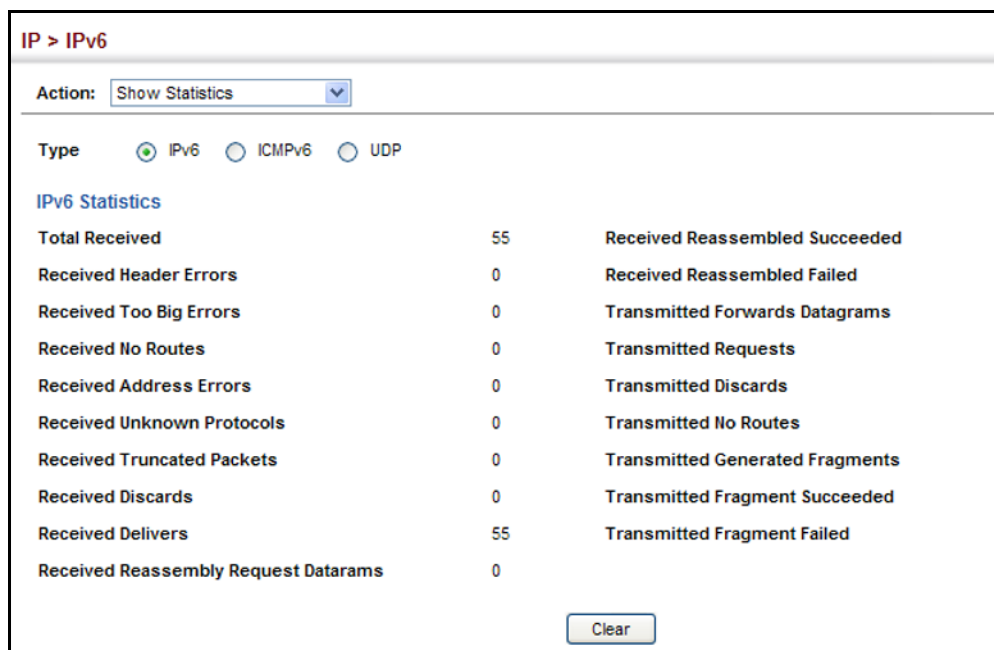


Figure 294: Showing IPv6 Statistics (ICMPv6)

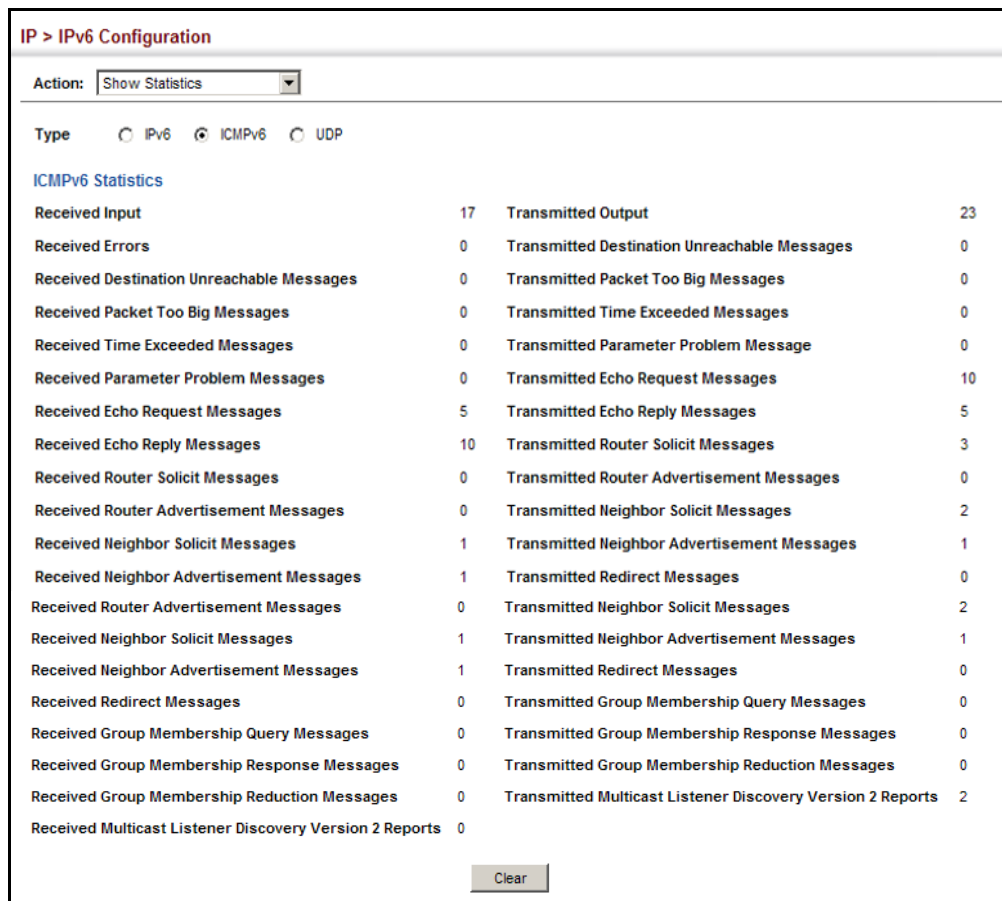
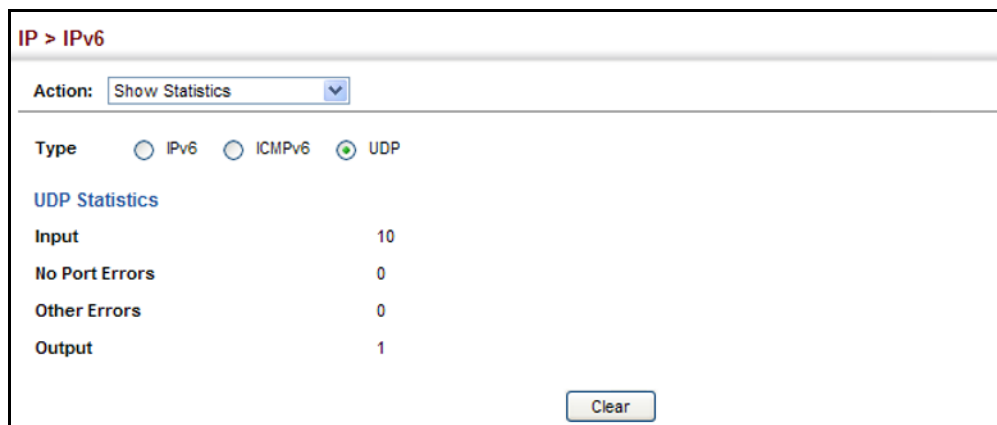


Figure 295: Showing IPv6 Statistics (UDP)



Showing the MTU for Responding Destinations Use the IP > IPv6 Configuration (Show MTU) page to display the maximum transmission unit (MTU) cache for destinations that have returned an ICMP packet-too-big message along with an acceptable MTU to this switch.

Parameters

These parameters are displayed:

Table 34: 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 296: Showing Reported MTU Values



This chapter describes how to configure Domain Name Service (DNS) on this switch. For information on DHCP snooping which is included in this folder, see [“DHCP Snooping” on page 346](#).

This chapter provides information on the following IP services, including:

- ◆ [DNS](#) – Configures default domain names, identifies servers to use for dynamic lookup, and shows how to configure static entries.
- ◆ [Multicast DNS](#) – Configures multicast DNS host name-to-address mapping on the local network without the need for a dedicated DNS server.
- ◆ [DHCP Client](#) – Specifies the DHCP client identifier for an interface.
- ◆ [DHCP Relay Option 82](#) – Enables DHCP relay service for attached host devices, including DHCP option 82 information.

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 456](#)).

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 297: Configuring General Settings for DNS

The screenshot shows a web interface for configuring DNS settings. At the top, the breadcrumb path is "IP Service > DNS > General". Below this, there is a section for "Action:" with a dropdown menu currently showing "Configure Global". Underneath, the "Domain Lookup" setting is checked with a green box and labeled "Enabled". The "Default Domain Name" is set to "my.site.com" in a text input field. At the bottom right of the configuration area, there are two buttons: "Apply" and "Revert".

Configuring a List of Domain Names Use the IP Service > DNS - General (Add Domain Name) page to configure a list of domain names to be tried in sequential order.

Command Usage

- ◆ Use this page to define a list of domain names that can be appended to incomplete host names (i.e., host names passed from a client that are not formatted with dotted notation).
- ◆ If there is no domain list, the default domain name is used (see [“Configuring General DNS Service Parameters” on page 453](#)). 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 456).

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 298: Configuring a List of Domain Names for DNS

The screenshot shows a web interface with a breadcrumb trail 'IP Service > DNS > General'. Below the breadcrumb, there is an 'Action:' dropdown menu set to 'Add Domain Name'. Underneath, there is a 'Domain Name' input field containing 'sample.com.uk'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

To show the list domain names:

1. Click IP Service, DNS.
2. Select Show Domain Names from the Action list.

Figure 299: Showing the List of Domain Names for DNS

The screenshot shows the same web interface as Figure 298, but the 'Action:' dropdown menu is now set to 'Show Domain Names'. Below the dropdown, there is a section titled 'Domain Name List' with a 'Total: 2' indicator. This section contains a table with two rows of domain names: 'google.com' and 'hinet.net'. Each row has a checkbox in the first column. At the bottom right, there are two buttons: 'Delete' and 'Revert'.

<input type="checkbox"/>	Domain Name
<input type="checkbox"/>	google.com
<input type="checkbox"/>	hinet.net

Configuring a List of Name Servers Use the IP Service > DNS - General (Add Name Server) page to configure a list of name servers to be tried in sequential order.

Command Usage

- ◆ To enable DNS service on this switch, configure one or more name servers, and enable domain lookup status (see [“Configuring General DNS Service Parameters” on page 453](#)).
- ◆ 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:

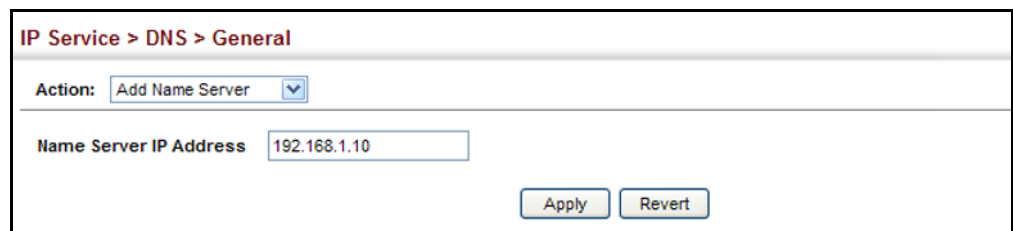
Name Server IP Address – Specifies the IPv4 or IPv6 address of a domain name server to use for name-to-address resolution. Up to six IP addresses can be added to the name server list.

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 300: Configuring a List of Name Servers for DNS

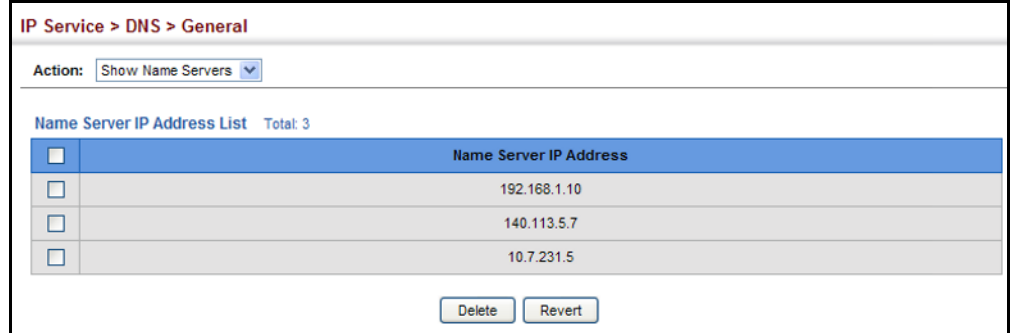


The screenshot shows a web interface for configuring DNS. At the top, the breadcrumb path is "IP Service > DNS > General". Below this, there is an "Action:" label followed by a dropdown menu currently set to "Add Name Server". Underneath, there is a "Name Server IP Address" label followed by a text input 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 301: 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:

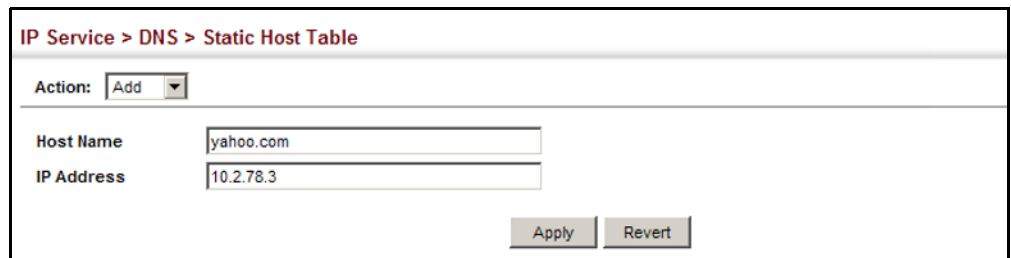
- ◆ **Host Name** – Name of a host device that is mapped to one or more IP addresses. (Range: 1-127 characters)
- ◆ **IP Address** – IPv4 or IPv6 address(es) associated with a host name.

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 302: Configuring Static Entries in the DNS Table



To show static entries in the DNS table:

1. Click IP Service, DNS, Static Host Table.
2. Select Show from the Action list.

Figure 303: Showing Static Entries in the DNS Table

<input type="checkbox"/>	Host	IP Address
<input type="checkbox"/>	yahoo.com	10.2.78.3
<input type="checkbox"/>	hinet.net	124.29.31.155
<input type="checkbox"/>	google.com	133.45.211.18

Displaying the DNS Cache Use the IP Service > DNS - Cache page to display entries in the DNS cache that have been learned via the designated name servers.

Command Usage

- ◆ Servers or other network devices may support one or more connections via multiple IP addresses. If more than one IP address is associated with a host name via information returned from a name server, a DNS client can try each address in succession, until it establishes a connection with the target device.

Parameters

These parameters are displayed:

- ◆ **No.** – The entry number for each resource record.
- ◆ **Flag** – The flag is always “4” indicating a cache entry and therefore unreliable.
- ◆ **Type** – This field includes CNAME which specifies the host address for the owner, and ALIAS which specifies an alias.
- ◆ **IP** – The IP address associated with this record.
- ◆ **TTL** – The time to live reported by the name server.
- ◆ **Host** – The host name associated with this record.

Web Interface

To display entries in the DNS cache:

1. Click IP Service, DNS, Cache.

Figure 304: Showing Entries in the DNS Cache

The screenshot shows a web interface titled "IP Service > DNS > Cache". Below the title, it says "Cache Information Total: 3". There is a table with the following data:

No.	Flag	Type	IP	TTL	Host
1	4	CNAME	192.168.110.2	360	www.sina.com.cn
2	4	CNAME	10.2.44.3	892	www.yahoo.akadns.new
3	4	ALIAS	pointer to: 2	298	www.yahoo.com

Below the table is a "Clear" button.

Multicast Domain Name Service

Use the IP Service > Multicast DNS page to enable multicast DNS host name-to-address mapping on the local network without the need for a dedicated DNS server.

Command Usage

- ◆ mDNS allows a network device to choose a domain name in the local DNS name space and announce it using a special multicast IP address. This allows any user to give their computers a link-local mDNS host name of the form "single-dns-label.local." Any name ending in ".local." is therefore link-local, and names within this domain are meaningful only on the link where they originate.
- ◆ When looking for the given host's IP address, the client sends a single-shot mDNS IP multicast query message to all the hosts sharing its local network. Any DNS query for a name ending with ".local." is sent to the mDNS multicast address 224.0.0.251 (or its IPv6 equivalent FF02::FB).

The corresponding host replies with a multicast message announcing itself. All machines in the subnet can then update their mDNS cache with the host's information sent in the reply message.

- ◆ To maintain an on-going cache of host names requires a process of continuous multicast DNS querying. This is done in several phases:
 - Probing – The DNS responder sends a probe message to the local network in order to verify that each entry its local cache is unique.
 - Announcing – The responder sends an unsolicited mDNS Response containing all of its newly registered resource records (both shared records, and unique records that have completed the probing step).
 - Updating – The responder repeats the Announcing step to update neighbor caches when the data for any local mDNS record changes.

Parameters

These parameters are displayed:

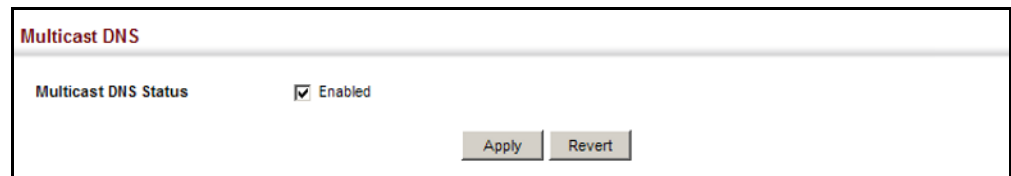
- ◆ **Multicast DNS Status** – Enables multicast DNS host name-to-address mapping on the local network. (Default: Enabled)

Web Interface

To configure multicast DNS:

1. Click IP Service, Multicast DNS.
2. Mark the check box to enable or disable mDNS as required
3. Click Apply.

Figure 305: Configuring Multicast DNS



Multicast DNS

Multicast DNS Status Enabled

Apply Revert

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.

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.
- ◆ **Text** – A text string. (Range: 1-32 characters)
- ◆ **Hex** – A hexadecimal value. (Range: 1-64 characters)

Web Interface

To configure a DHCP client identifier:

1. Click IP Service, DHCP, Client.
2. Mark the check box to enable this feature. Select the default setting, or the format for a vendor class identifier. If a non-default value is used, enter a text string or hexadecimal value.
3. Click Apply.

Figure 306: Specifying A DHCP Client Identifier

IP Service > DHCP > Client

VLAN: 1

Vendor Class ID: Default ECS4210-12P

Apply Revert

Configuring DHCP Relay Option 82

Use the IP Service > DHCP > Relay Option 82 page to configure DHCP relay service for attached host devices, including DHCP option 82 information. DHCP provides an option for sending information about its DHCP clients to the DHCP server (specifically, the interface on the relay server through which the DHCP client request was received). Also known as DHCP Relay Option 82, it allows compatible DHCP servers to use this information when assigning IP addresses, or to set other services or policies for clients.

Option 82 information contains information which can identify both the relay agent and the interface through which the DHCP request was received:

- ◆ The DHCP Relay Information Option Remote ID (RID) is the access node identifier – a string used to identify the switch to the DHCP server.


- ◆ The DHCP Relay Information Option Fields are the Option 82 circuit identification fields (CID – including VLAN ID, stack unit, and port). These fields identify the requesting device by indicating the interface through which the relay agent received the request.

If DHCP relay is enabled, and this switch sees a DHCP client request, it inserts its own IP address into the request so that the DHCP server will know the subnet where the client is located. Depending on the selected frame format set for the remote-id, this information may specify the MAC address, IP address, or an arbitrary string for the requesting device (that is, the relay agent in this context).

By default, the relay agent also fills in the Option 82 circuit-id field with information indicating the local interface over which the switch received the DHCP client request, including the VLAN ID, stack unit, and port. This allows DHCP client-server exchange messages to be forwarded between the server and client without having to flood them onto the entire VLAN.

The switch then forwards the packet to the DHCP server. When the server receives the DHCP request, it allocates a free IP address for the DHCP client from its defined scope for the DHCP client's subnet, and sends a DHCP response back to the DHCP relay agent (i.e., this switch). This switch then passes the DHCP response received from the server to the client.

Figure 307: Layer 2 DHCP Relay Service



Command Usage

- ◆ You must specify the IP address for at least one active DHCP server. Otherwise, the switch's DHCP relay agent will not be able to forward client requests to a DHCP server. Up to five DHCP servers can be specified in order of preference.

If any of the specified DHCP server addresses are not located in the same network segment with this switch, specify the default router through which this switch can reach other IP subnetworks (see ["Setting the Switch's IP Address \(IP Version 4\)" on page 431](#) or ["Configuring the IPv6 Default Gateway" on page 434](#)).

- ◆ DHCP Snooping Information Option 82 (see [page 348](#)) and DHCP Relay Information Option 82 cannot both be enabled at the same time.
- ◆ DHCP request packets received by the switch are handled as follows:
 - If a DHCP relay server has been set on the switch, when the switch receives a DHCP request packet *without* option 82 information from the

management VLAN or a non-management VLAN, it will add option 82 relay information and the relay agent's address to the DHCP request packet, and then unicast it to the DHCP server.

- If a DHCP relay server has been set on the switch, when the switch receives a DHCP request packet *with* option 82 information from the management VLAN or a non-management VLAN, it will process it according to the configured relay information option policy:
 - If the policy is "replace," the DHCP request packet's option 82 content (the RID and CID sub-option) is replaced with information provided by the switch. The relay agent address is inserted into the DHCP request packet, and the switch then unicasts this packet to the DHCP server.
 - If the policy is "keep," the DHCP request packet's option 82 content will be retained. The relay agent address is inserted into the DHCP request packet, and the switch then unicasts this packet to the DHCP server.
 - If the policy is "drop," the original DHCP request packet is flooded onto the VLAN which received the packet but is not relayed.

- ◆ DHCP reply packets received by the relay agent are handled as follows:

When the relay agent receives a DHCP reply packet with Option 82 information over the management VLAN, it first ensures that the packet is destined for it.

- If the RID in the DHCP reply packet is not identical with that configured on the switch, the option 82 information is retained, and the packet is flooded onto the VLAN through which it was received.
- If the RID in the DHCP reply packet matches that configured on the switch, it then removes the Option 82 information from the packet, and sends it on as follows:
 - If the DHCP packet's broadcast flag is on, the switch uses the circuit-id information contained in the option 82 information fields to identify the VLAN connected to the requesting client and then broadcasts the DHCP reply packet to this VLAN.
 - If the DHCP packet's broadcast flag is off, the switch uses the circuit-id information in option 82 fields to identify the interface connected to the requesting client and unicasts the reply packet to the client.
- ◆ DHCP packets are flooded onto the VLAN which received them if DHCP relay service is enabled on the switch and any of the following situations apply:
 - There is no DHCP relay server set on the switch, when the switch receives a DHCP packet.

- A DHCP relay server has been set on the switch, when the switch receives a DHCP request packet with a non-zero relay agent address field (that is not the address of this switch).
- A DHCP relay server has been set on the switch, when the switch receives DHCP reply packet without option 82 information from the management VLAN.
- The reply packet contains a valid relay agent address field (that is not the address of this switch), or receives a reply packet with a zero relay agent address through the management VLAN.
- A DHCP relay server has been set on the switch, and the switch receives a reply packet on a non-management VLAN.

Parameters

These parameters are displayed:

- ◆ **Insertion of Relay Information** – Enable DHCP Option 82 information relay. (Default: Disabled)
- ◆ **DHCP Option Policy** – Specifies how to handle client requests which already contain DHCP Option 82 information:
 - **Drop** - Floods the original request packet onto the VLAN that received it instead of relaying it. (This is the default.)
 - **Keep** - Retains the Option 82 information in the client request, inserts the relay agent's address, and unicasts the packet to the DHCP server.
 - **Replace** - Replaces the Option 82 information circuit-id and remote-id fields in the client's request with information provided by the relay agent itself, inserts the relay agent's address, and unicasts the packet to the DHCP server.
- ◆ **DHCP Sub-option Format** – Specifies whether or not to use the sub-type and sub-length fields in the circuit-ID (CID) and remote-ID (RID) in Option 82 information. (Default: Included)
- ◆ **DHCP Remote ID Sub-option** – Specifies the format used to identify this switch as the DHCP relay agent to the DHCP server:
 - **MAC Address** - Includes a MAC address field for the relay agent (that is, the MAC address of the switch's CPU). This attribute can be encoded in hexadecimal or ASCII.
 - **IP Address** - Includes the IP address field for the relay agent (that is, 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)

- ◆ **Server IP Address** – Addresses of DHCP servers or relay servers to be used by the switch's DHCP relay agent in order of preference.

Web Interface

To configure DHCP relay service:

1. Click IP Service, DHCP, Relay Option 82.
2. Enable or disable Option 82.
3. Set the Option 82 policy to specify how to handle Option 82 information already contained in DHCP client request packets.
4. Specify whether or not include "type" and "length" sub-options.
5. Set the frame format used for the remote ID.
6. Enter up to five IP addresses for DHCP servers or relay servers in order of preference.
7. Click Apply.

Figure 308: Configuring DHCP Relay Information Option 82 Service

IP Service > DHCP > Relay Option 82

Information

Insertion of Relay Information Enabled

DHCP Option Policy

DHCP Sub-option Format

DHCP Remote ID Sub-option

Server

Server IP Address

Note: DHCP relay will not be active if an active DHCP server is detected on the same network segment.

Multicast Filtering


This chapter describes how to configure the following multicast services:

- ◆ **IGMP** – Configuring snooping and query parameters.
- ◆ **Filtering and Throttling** – Filtering specified multicast service, or throttling the maximum of multicast groups allowed on an interface.
- ◆ **MLD Snooping** – Configures snooping and query parameters for IPv6.
- ◆ **Multicast VLAN Registration (MVR)** – 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 309: 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” on page 502](#).

Layer 2 IGMP (Snooping and Query)

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 470](#)) 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 refused¹⁰. The switch maintains information about multicast groups, 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). For IGMPv1/v2/v3 hosts, the source address of a channel is always null (indicating that any source is acceptable).

10. Source IP lists is not supported in IGMPv3 reports by the switch due to an ASIC limitation.



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 473](#)). 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 255 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 470](#)).

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 473](#)). 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 476](#)).

IGMP Snooping with Proxy Reporting – The switch supports last leave, and query suppression (as defined in DSL Forum TR-101, April 2006):

- ◆ When proxy reporting is disabled, all IGMP reports received by the switch are forwarded natively to the upstream multicast routers.
- ◆ Last Leave: Intercepts, absorbs and summarizes IGMP leaves coming from IGMP hosts. IGMP leaves are relayed upstream only when necessary, that is, when the last user leaves a multicast group.
- ◆ Query Suppression: Intercepts and processes IGMP queries in such a way that IGMP specific queries are never sent to client ports.

The only deviation from TR-101 is that the marking of IGMP traffic initiated by the switch with priority bits as defined in R-250 is not supported.

Configuring IGMP Snooping and Query Parameters

Use the Multicast > IGMP Snooping > General page to configure the switch to forward multicast traffic intelligently. Based on the IGMP query and report messages, the switch forwards multicast traffic only to the ports that request it. This prevents the switch from broadcasting the traffic to all ports and possibly disrupting network performance.

Command Usage

- ◆ **IGMP Snooping** – This switch can passively snoop on IGMP Query and Report packets transferred between IP multicast routers/switches and IP multicast host groups to identify the IP multicast group members. It simply monitors the IGMP packets passing through it, picks out the group registration information, and configures the multicast filters accordingly.



Note: If unknown multicast traffic enters a VLAN which has been configured with a router port, the traffic is forwarded to that port. However, if no router port exists on the VLAN, the traffic is dropped if unregistered data flooding is disabled (default behavior), or flooded throughout the VLAN if unregistered data flooding is enabled (see “Unregistered Data Flooding” in the Command Attributes section).

- ◆ **IGMP Querier** – A router, or multicast-enabled switch, can periodically ask their hosts if they want to receive multicast traffic. If there is more than one router/switch on the LAN performing IP multicasting, one of these devices is elected “querier” and assumes the role of querying the LAN for group members. It then propagates the service requests on to any upstream multicast switch/router to ensure that it will continue to receive the multicast service.



Note: Multicast routers use this information from IGMP snooping and query reports, along with a multicast routing protocol such as DVMRP or PIM, to support IP multicasting across the Internet.

Parameters

These parameters are displayed:

- ◆ **IGMP Snooping Status** – When enabled, the switch will monitor network traffic to determine which hosts want to receive multicast traffic. This is referred to as IGMP Snooping. (Default: Disabled)

When IGMP snooping is enabled globally, the per VLAN interface settings for IGMP snooping take precedence (see “[Setting IGMP Snooping Status per Interface](#)” on page 478).

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 or group-and-source-specific queries, each with a large source list and the Maximum Response Time set to a large value.

To protect against this kind of attack, (1) routers should not forward queries. This is easier to accomplish if the query carries the Router Alert option. (2) Also, when the switch is acting in the role of a multicast host (such as when using proxy routing), it should ignore version 2 or 3 queries that do not contain the Router Alert option.

- ◆ **Unregistered Data Flooding** – Floods unregistered multicast traffic into the attached VLAN. (Default: Disabled)

Once the table used to store multicast entries for IGMP snooping and multicast routing is filled, no new entries are learned. If no router port is configured in the attached VLAN, and unregistered-flooding is disabled, any subsequent multicast traffic not found in the table is dropped, otherwise it is flooded throughout the VLAN.

- ◆ **Version Exclusive** – Discards any received IGMP messages which use a version different to that currently configured by the IGMP Version attribute. (Default: Disabled)

- ◆ **IGMP Unsolicited Report Interval** – Specifies how often the upstream interface should transmit unsolicited IGMP reports when proxy reporting is enabled. (Range: 1-65535 seconds, Default: 400 seconds)

When a new upstream interface (that is, uplink port) starts up, the switch sends unsolicited reports for all currently learned multicast channels via the new upstream interface.

This command only applies when proxy reporting is enabled.

- ◆ **Router Port Expire Time** – The time the switch waits after the previous querier stops before it considers it to have expired. (Range: 1-65535, Recommended Range: 300-500 seconds, Default: 300)

- ◆ **IGMP Snooping Version** – Sets the protocol version for compatibility with other devices on the network. This is the IGMP Version the switch uses to send snooping reports. (Range: 1-3; Default: 2)

This attribute configures the IGMP report/query version used by IGMP snooping. Versions 1 - 3 are all supported, and versions 2 and 3 are backward compatible, so the switch can operate with other devices, regardless of the snooping version employed.

- ◆ **Querier Status** – When enabled, the switch can serve as the Querier, which is responsible for asking hosts if they want to receive multicast traffic. This feature is not supported for IGMPv3 snooping. (Default: Disabled)

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 310: Configuring General Settings for IGMP Snooping

Multicast > IGMP Snooping > General	
IGMP Snooping Status	<input checked="" type="checkbox"/> Enabled
Proxy Reporting Status	<input type="checkbox"/> Enabled
TCN Flood	<input type="checkbox"/> Enabled
TCN Query Solicit	<input type="checkbox"/> Enabled
Router Alert Option	<input type="checkbox"/> Enabled
Unregistered Data Flooding	<input type="checkbox"/> Enabled
Version Exclusive	<input type="checkbox"/> Enabled
IGMP Unsolicited Report Interval (1-65535)	<input type="text" value="400"/> seconds
Router Port Expire Time (1-65535)	<input type="text" value="300"/> seconds
IGMP Snooping Version (1-3)	<input type="text" value="2"/>
Querier Status	<input type="checkbox"/> Enabled

Specifying Static Interfaces for 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 470](#)) 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.

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 311: Configuring a Static Interface for a Multicast Router

The screenshot shows the configuration page for a Multicast Router. The breadcrumb is 'Multicast > IGMP Snooping > Multicast Router'. The 'Action' dropdown is set to 'Add Static Multicast Router'. Below this, there is a 'VLAN' dropdown set to '1'. Under the 'Interface' section, the 'Port' radio button is selected and set to '1', while the 'Trunk' radio button is unselected. At the bottom right, there are 'Apply' and 'Revert' buttons.

To show the static interfaces attached to a multicast router:

1. Click Multicast, 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 312: Showing Static Interfaces Attached a Multicast Router

The screenshot shows the configuration page for a Multicast Router. The breadcrumb is 'Multicast > IGMP Snooping > Multicast Router'. The 'Action' dropdown is set to 'Show Static Multicast Router'. Below this, there is a 'VLAN' dropdown set to '1'. The main content area is titled 'Static Multicast Router Interface List' with a 'Total: 6' indicator. It contains a table with the following data:

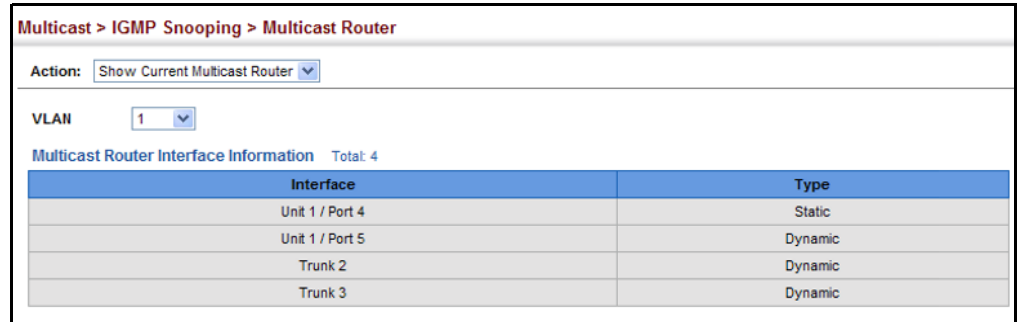
<input type="checkbox"/>	Interface
<input type="checkbox"/>	Unit 1 / Port 1
<input type="checkbox"/>	Unit 1 / Port 2
<input type="checkbox"/>	Unit 1 / Port 3
<input type="checkbox"/>	Trunk 2
<input type="checkbox"/>	Trunk 5
<input type="checkbox"/>	Unit 1 / Port 4

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

To show the all interfaces attached to a multicast router:

1. Click Multicast, IGMP Snooping, Multicast Router.
2. Select Current Multicast Router from the Action list.
3. Select the VLAN for which to display this information. Ports in the selected VLAN which are attached to a neighboring multicast router/switch are displayed.

Figure 313: Showing Current Interfaces Attached a Multicast Router



Multicast > IGMP 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 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 470](#)). 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 314: Assigning an Interface to a Multicast Service

Multicast > IGMP Snooping > IGMP Member

Action: Add Static Member

VLAN: 1

Interface: Port 1 Trunk 1

Multicast IP: 224.1.1.1

Apply 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 315: Showing Static Interfaces Assigned to a Multicast Service

Multicast > IGMP Snooping > IGMP Member

Action: Show Static Member

VLAN: 1

IGMP Member Interface List Total: 6

	Interface	Multicast IP
<input type="checkbox"/>	Unit 1 / Port 1	224.1.1.1
<input type="checkbox"/>	Unit 1 / Port 2	224.1.2.2
<input type="checkbox"/>	Unit 1 / Port 3	230.1.1.1
<input type="checkbox"/>	Trunk 2	230.1.2.2
<input type="checkbox"/>	Trunk 5	239.1.1.1
<input type="checkbox"/>	Unit 1 / Port 4	239.2.2.2

Delete Revert

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 470.](#)

Command Usage

Multicast Router Discovery

There have been many mechanisms used in the past to identify multicast routers. This has led to interoperability issues between multicast routers and snooping switches from different vendors. In response to this problem, the Multicast Router Discovery (MRD) protocol has been developed for use by IGMP snooping and multicast routing devices. MRD is used to discover which interfaces are attached to multicast routers, allowing IGMP-enabled devices to determine where to send multicast source and group membership messages. (MRD is specified in draft-ietf-magma-mrdisc-07.)

Multicast source data and group membership reports must be received by all multicast routers on a segment. Using the group membership protocol query messages to discover multicast routers is insufficient due to query suppression. MRD therefore provides a standardized way to identify multicast routers without relying on any particular multicast routing protocol.



Note: The default values recommended in the MRD draft are implemented in the switch.

Multicast Router Discovery uses the following three message types to discover multicast routers:

- ◆ Multicast Router Advertisement – Advertisements are sent by routers to advertise that IP multicast forwarding is enabled. These messages are sent unsolicited periodically on all router interfaces on which multicast forwarding is enabled. They are sent upon the occurrence of these events:
 - Upon the expiration of a periodic (randomized) timer.
 - As a part of a router's start up procedure.
 - During the restart of a multicast forwarding interface.
 - On receipt of a Solicitation message.
- ◆ Multicast Router Solicitation – Devices send Solicitation messages in order to solicit Advertisement messages from multicast routers. These messages are used to discover multicast routers on a directly attached link. Solicitation messages are also sent whenever a multicast forwarding interface is initialized or re-initialized. Upon receiving a solicitation on an interface with IP multicast forwarding and MRD enabled, a router will respond with an Advertisement.

- ◆ Multicast Router Termination – These messages are sent when a router stops IP multicast routing functions on an interface. Termination messages are sent by multicast routers when:
 - Multicast forwarding is disabled on an interface.
 - An interface is administratively disabled.
 - The router is gracefully shut down.

Advertisement and Termination messages are sent to the All-Snoopers multicast address. Solicitation messages are sent to the All-Routers multicast address.



Note: MRD messages are flooded to all ports in a VLAN where IGMP snooping or routing has been enabled. To ensure that older switches which do not support MRD can also learn the multicast router port, the switch floods IGMP general query packets, which do not have a null source address (0.0.0.0), to all ports in the attached VLAN. IGMP packets with a null source address are only flooded to all ports in the VLAN if the system is operating in multicast flooding mode, such as when a new VLAN or new router port is being established, or a spanning tree topology change has occurred. Otherwise, this kind of packet is only forwarded to known multicast routing ports.

Parameters

These parameters are displayed:

- ◆ **VLAN** – ID of configured VLANs. (Range: 1-4094)
- ◆ **IGMP Snooping Status** – When enabled, the switch will monitor network traffic on the indicated VLAN interface to determine which hosts want to receive multicast traffic. This is referred to as IGMP Snooping. (Default: Disabled)

When IGMP snooping is enabled globally (see [page 470](#)), 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. (Default: Disabled)

If version exclusive is disabled on a VLAN, then this setting is based on the global setting configured on the Multicast > IGMP Snooping > General page. If it is enabled on a VLAN, then this setting takes precedence over the global setting.
- ◆ **Immediate Leave Status** – Immediately deletes a member port of a multicast service if a leave packet is received at that port and immediate leave is enabled for the parent VLAN. (Default: Disabled)

If immediate leave is not used, a multicast router (or querier) will send a group-specific query message when an IGMPv2 group leave message is received. The router/querier stops forwarding traffic for that group only if no host replies to the query within the specified time out period. Note that this time out is set to Last Member Query Interval * Robustness Variable (fixed at 2) as defined in RFC 2236.

If immediate leave is enabled, the switch assumes that only one host is connected to the interface. Therefore, immediate leave should only be enabled on an interface if it is connected to only one IGMP-enabled device, either a service host or a neighbor running IGMP snooping.

This attribute is only effective if IGMP snooping is enabled, and IGMPv2 snooping is used.

- ◆ **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. (Default: Based on global setting)

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. (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.

- ◆ **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 470](#)), or as a proxy host when IGMP snooping proxy reporting is enabled ([page 470](#)).

- ◆ **Query Response Interval** – The maximum time the system waits for a response to general queries. (Range: 10-31744 tenths of a second; Default: 10 seconds)

This attribute applies when the switch is serving as the querier ([page 470](#)), or as a proxy host when IGMP snooping proxy reporting is enabled ([page 470](#)).

- ◆ **Last Member Query Interval** – The interval to wait for a response to a group-specific or group-and-source-specific query message. (Range: 1-31740 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 query message, and starts a timer. If no reports are received before the timer expires, the group record is deleted, and a report is sent to the upstream multicast router.

A reduced value will result in reduced time to detect the loss of the last member of a group or source, but may generate more burst traffic.

This attribute will take effect only if IGMP snooping proxy reporting is enabled (see [page 470](#)) or IGMP querier is enabled ([page 470](#)).

- ◆ **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 316: Configuring IGMP Snooping on an Interface

The screenshot shows a web interface titled "Multicast > IGMP Snooping > Interface". At the top, there is a dropdown menu for "Action:" set to "Configure VLAN". Below this, a table of configuration options is displayed:

VLAN	1
IGMP Snooping Status	<input checked="" type="checkbox"/> Enabled
Version Exclusive	<input type="checkbox"/> Enabled
Immediate Leave Status	<input type="checkbox"/> Enabled
Multicast Router Discovery	<input type="checkbox"/> Enabled
General Query Suppression	<input type="checkbox"/> Enabled
Proxy Reporting	Disabled
Interface Version (1-3)	2
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

At the bottom right of the form, there are two buttons: "Apply" and "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 317: Showing Interface Settings for IGMP Snooping

VLAN	IGMP Snooping Status	Immediate Leave Status	Query Interval	Query Response Interval	Last Member Query Interval	Last Member Query Count	Proxy (Query) Address	Proxy Reporting	Multicast Router Discovery	General Query Suppression	Version Exclusive	Interface Version
1	Enabled	Disabled	10	100	10	2	10.1.1.1	Enabled	Enabled	Disabled	Enabled	1
2	Disabled	Disabled	10	100	10	2	20.2.2.2	Disabled	Disabled	Enabled	Disabled	3
3	Disabled	Disabled	10	100	10	2	30.3.3.3	Disabled	Enabled	Disabled	Disabled	2
10	Disabled	Disabled	10	100	10	2	100.10.10.10	Disabled	Disabled	Enabled	Disabled	1

Filtering Multicast Data at Interfaces

Use the Multicast > IGMP Snooping > Interface (Configure Port/Trunk) page to configure an interface to drop IGMP query packets.

Parameters

These parameters are displayed:

- ◆ **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. (Default: Disabled)
- ◆ **IGMP RADIUS Authentication** – Sends access request to RADIUS server for authentication when a join report is received on an interface. (Default: Disabled)

If the RADIUS server responds with authentication success, the switch learns the group report and will not send the access request to the RADIUS server when receiving the same report again in 86400 seconds (1 day).

When receiving an IGMPv3 report message, the switch will send the access request to the RADIUS server only when the record type is IS_EX (MODE_IS_EXCLUDE) which excludes a source list or TO_EX (CHANGE_TO_EXCLUDE_MODE), and the source list is empty. Other type of packets will not be authenticated.

The first time the report is received and is being authenticated, regardless of whether the authentication result succeeds or fails, this report will still be sent to any configured multicast router port.

The attribute value pairs configured on the RADIUS server are shown below.

Table 35: RADIUS Server AVPs

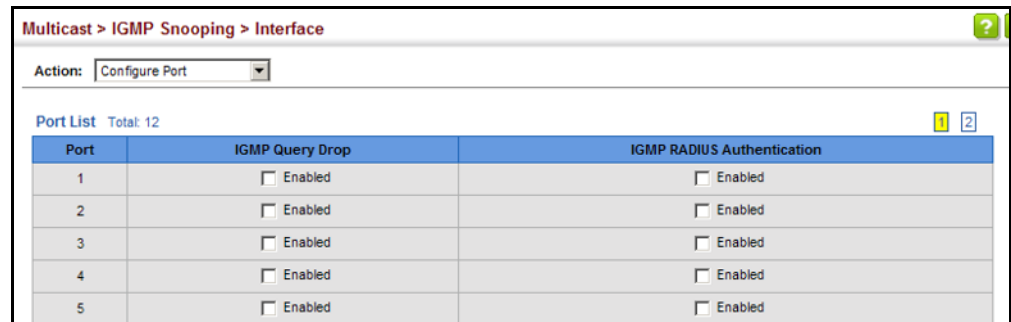
Attribute Name	AVP Type	Entry
USER_NAME	1	User MAC address
USER_PASSWORD	2	User MAC address
NAS_IP_ADDRESS	4	Switch IP
NAS_PORT	5	User port
FRAMED_IP_ADDRESS	8	Multicast group IP

Web Interface

To drop IGMP query packets or multicast data packets:

1. Click Multicast, IGMP Snooping, Interface.
2. Select Configure Port or Configure Trunk from the Action List.
3. Enable the required drop functions for any interface.
4. Click Apply.

Figure 318: Dropping IGMP Query Packets



Displaying Multicast Groups Discovered by IGMP Snooping

Use the Multicast > IGMP Snooping > Forwarding Entry page to display the forwarding entries learned through IGMP Snooping.

Command Usage

To display information about multicast groups, IGMP Snooping must first be enabled on the switch (see [page 470](#)).

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 319: Showing Multicast Groups Learned by IGMP Snooping

Multicast > IGMP Snooping > Forwarding Entry

VLAN: 1

IGMP Snooping Forwarding Entry List Total: 10

Group Address	Interface	Up Time	Expire	Count
224.1.1.1	Eth 1 / 9 (Router Port)	00:00:06:46		2 (Port)
	Eth 1 / 11 (Member Port)	00:00:06:46	03:46	1 (Host)
224.1.1.2	Eth 1 / 9 (Router Port)		02:24	1 (Port)
	Eth 1 / 9 (Router Port)	00:00:16:14		1 (Port)
239.255.255.250	Eth 1 / 9 (Router Port)	00:00:08:47		2 (Port)
	Eth 1 / 11 (Member Port)	00:00:08:47	03:46	1 (Host)

Clear Click this button to clear all IGMP Snooping dynamic groups.

Displaying IGMP Snooping Statistics Use the Multicast > IGMP Snooping > Statistics pages to display IGMP snooping protocol-related statistics for the specified interface.

Parameters

These parameters are displayed:

- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Port** – Port identifier. (Range: 1-12/28)

- ◆ **Trunk** – Trunk identifier. (Range: 1-8/12)

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 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.

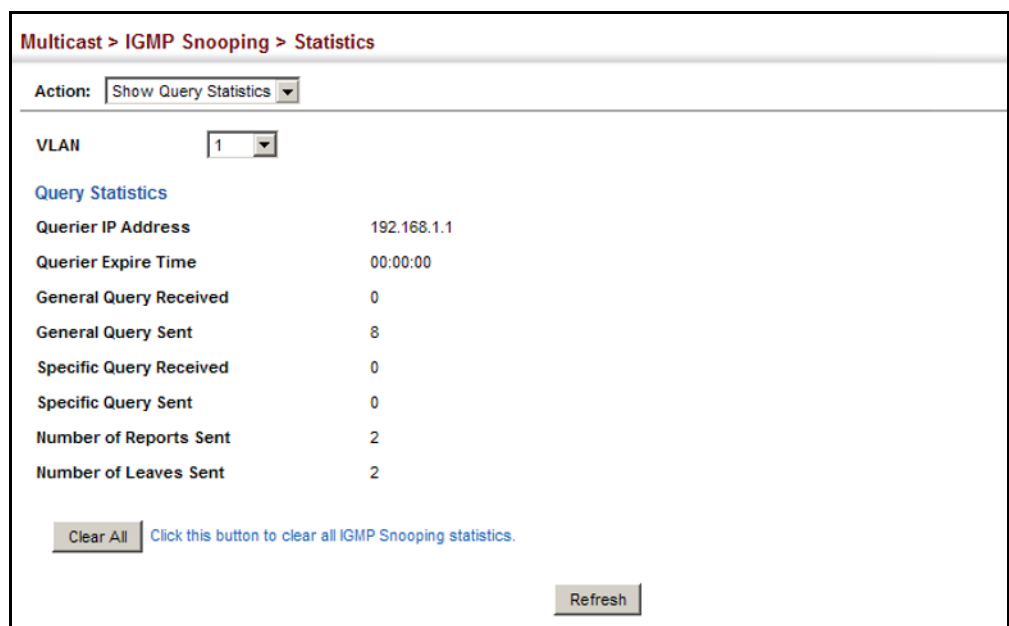
- ◆ **G(-S)-S Query** – The number of group specific or group-and-source specific query messages sent from this interface.

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 320: Displaying IGMP Snooping Statistics – Query



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 321: Displaying IGMP Snooping Statistics – VLAN



To display IGMP snooping protocol-related statistics for a port:

1. Click Multicast, IGMP Snooping, Statistics.
2. Select Show Port Statistics from the Action list.
3. Select a Port.

Figure 322: Displaying IGMP Snooping Statistics – Port



Filtering and Throttling IGMP Groups

In certain switch applications, the administrator may want to control the multicast services that are available to end users. For example, an IP/TV service based on a specific subscription plan. The IGMP filtering feature fulfills this requirement by restricting access to specified multicast services on a switch port, and IGMP throttling limits the number of simultaneous multicast groups a port can join.

IGMP filtering enables you to assign a profile to a switch port that specifies multicast groups that are permitted or denied on the port. An IGMP filter profile can contain one or more addresses, or a range of multicast addresses; but only one profile can be assigned to a port. When enabled, IGMP join reports received on the port are checked against the filter profile. If a requested multicast group is permitted, the IGMP join report is forwarded as normal. If a requested multicast group is denied, the IGMP join report is dropped.

IGMP throttling sets a maximum number of multicast groups that a port can join at the same time. When the maximum number of groups is reached on a port, the switch can take one of two actions; either “deny” or “replace.” If the action is set to deny, any new IGMP join reports will be dropped. If the action is set to replace, the switch randomly removes an existing group and replaces it with the new multicast group.

Enabling IGMP Filtering and Throttling

Use the Multicast > IGMP Snooping > Filter (Configure General) page to enable IGMP filtering and throttling globally on the switch.

Parameters

These parameters are displayed:


- ◆ **IGMP Filter Status** – Enables IGMP filtering and throttling globally for the switch. (Default: Disabled)

Web Interface

To enable IGMP filtering and throttling on the switch:

1. Click Multicast, IGMP Snooping, Filter.
2. Select Configure General from the Step list.
3. Enable IGMP Filter Status.
4. Click Apply.

Figure 323: Enabling IGMP Filtering and Throttling



Configuring IGMP Filter Profiles

Use the Multicast > IGMP Snooping > Filter (Configure Profile – Add) page to create an IGMP profile and set its access mode. Then use the (Add Multicast Group Range) page to configure the multicast groups to filter.

Command Usage

Specify a range of multicast groups by entering a start and end IP address; or specify a single multicast group by entering the same IP address for the start and end of the range.

Parameters

These parameters are displayed:

Add

- ◆ **Profile ID** – Creates an IGMP profile. (Range: 1-4294967295)
- ◆ **Access Mode** – Sets the access mode of the profile; either permit or deny. (Default: Deny)

When the access mode is set to permit, IGMP join reports are processed when a multicast group falls within the controlled range. When the access mode is set to deny, IGMP join reports are only processed when the multicast group is not in the controlled range.

Add Multicast Group Range

- ◆ **Profile ID** – Selects an IGMP profile to configure.
- ◆ **Start Multicast IP Address** – Specifies the starting address of a range of multicast groups.
- ◆ **End Multicast IP Address** – Specifies the ending address of a range of multicast groups.

Web Interface

To create an IGMP filter profile and set its access mode:

1. Click Multicast, IGMP Snooping, Filter.
2. Select Configure Profile from the Step list.

3. Select Add from the Action list.
4. Enter the number for a profile, and set its access mode.
5. Click Apply.

Figure 324: Creating an IGMP Filtering Profile

Multicast > IGMP Snooping > Filter

Step: 2. Configure Profile Action: Add

Profile ID (1-4294967295) 19

Access Mode Permit

Apply Revert

To show the IGMP filter profiles:

1. Click Multicast, IGMP Snooping, Filter.
2. Select Configure Profile from the Step list.
3. Select Show from the Action list.

Figure 325: Showing the IGMP Filtering Profiles Created

Multicast > IGMP Snooping > Filter

Step: 2. Configure Profile Action: Show

IGMP Snooping Filter Profile List Total: 1

<input type="checkbox"/>	Profile ID	Action Mode
<input type="checkbox"/>	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 326: Adding Multicast Groups to an IGMP Filtering Profile

The screenshot shows the configuration page for adding multicast groups. The breadcrumb is 'Multicast > IGMP Snooping > Filter'. The 'Step' dropdown is set to '2. Configure Profile' and the 'Action' dropdown is set to 'Add Multicast Group Range'. The 'Profile ID' is set to '19'. The 'Start Multicast IP Address' is '239.2.3.1' and the 'End Multicast IP Address' is '239.2.3.200'. There are 'Apply' and 'Revert' buttons at the bottom right.

To show the multicast groups configured for an IGMP filter profile:

1. Click Multicast, IGMP Snooping, Filter.
2. Select Configure Profile from the Step list.
3. Select Show Multicast Group Range from the Action list.
4. Select the profile for which to display this information.

Figure 327: Showing the Groups Assigned to an IGMP Filtering Profile

The screenshot shows the configuration page for showing multicast groups. The breadcrumb is 'Multicast > IGMP Snooping > Filter'. The 'Step' dropdown is set to '2. Configure Profile' and the 'Action' dropdown is set to 'Show Multicast Group Range'. The 'Profile ID' is set to '19'. Below the form is a table titled 'Multicast IP Address Range List' with a 'Total: 1' indicator. The table has three columns: a checkbox, 'Start Multicast IP Address', and 'End Multicast IP Address'. The first row shows a checked checkbox, '239.2.3.1', and '239.2.3.200'. There are 'Delete' and 'Revert' buttons at the bottom right.

<input type="checkbox"/>	Start Multicast IP Address	End Multicast IP Address
<input checked="" type="checkbox"/>	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: 1-1023; Default: 1023)
- ◆ **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 328: Configuring IGMP Filtering and Throttling Interface Settings

Port	Profile ID	Max Multicast Groups (1-1023)	Current Multicast Groups	Throttling Action Mode	Throttling Status
1	19	64	1	Deny	False
2	(none)	1023	0	Deny	False
3	(none)	1023	0	Deny	False
4	(none)	1023	0	Deny	False
5	(none)	1023	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 329: Configuring General Settings for MLD Snooping

The screenshot shows the configuration page for MLD Snooping General settings. The breadcrumb trail is "Multicast > MLD Snooping > General". The settings are as follows:

MLD Snooping Status	<input type="checkbox"/> Enabled
Querier Status	<input type="checkbox"/> Enabled
Robustness (2-10)	<input type="text" value="2"/>
Query Interval (60-125)	<input type="text" value="125"/> seconds
Query Max Response Time (5-25)	<input type="text" value="10"/> seconds
Router Port Expiry Time (300-500)	<input type="text" value="300"/> seconds
MLD Snooping Version (1-2)	<input type="text" value="2"/>
Unknown Multicast Mode	<input type="text" value="To Router Port"/>

At the bottom right, there are two buttons: "Apply" and "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 330: 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 494) before a multicast router port can take effect.

Parameters

These parameters are displayed:

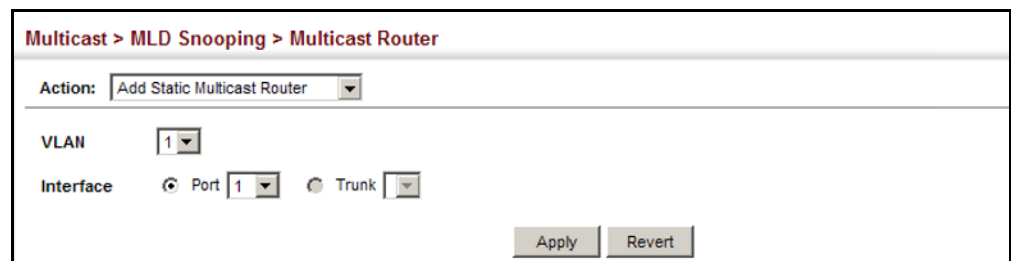
- ◆ **VLAN** – Selects the VLAN which is to propagate all IPv6 multicast traffic coming from the attached multicast router. (Range: 1-4094)
- ◆ **Interface** – Activates the Port or Trunk scroll down list.
- ◆ **Port or Trunk** – Specifies the interface attached to a multicast router.

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 331: Configuring a Static Interface for an IPv6 Multicast Router

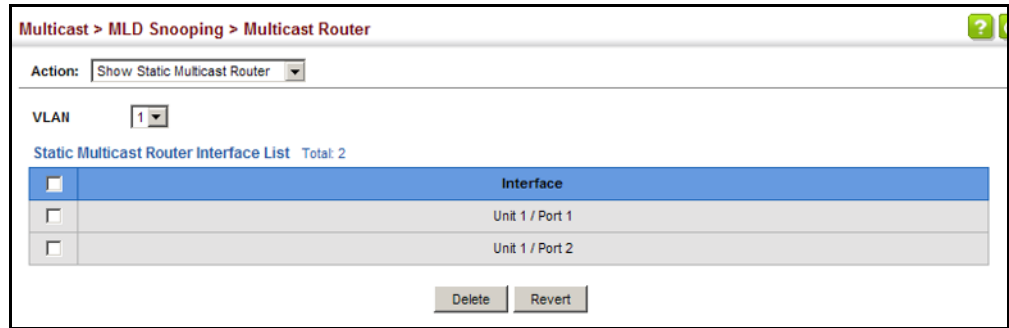


The screenshot shows a web interface for configuring a static interface for an IPv6 multicast router. The breadcrumb navigation is "Multicast > MLD Snooping > Multicast Router". The "Action" dropdown menu is set to "Add Static Multicast Router". Below this, there are two rows of configuration options. The first row is "VLAN" with a dropdown menu showing "1". The second row is "Interface" with two radio buttons: "Port" (selected) and "Trunk". Both radio buttons have a dropdown menu showing "1". At the bottom right of the form, there are two buttons: "Apply" and "Revert".

To show the static interfaces attached to a multicast router:

1. Click Multicast, MLD Snooping, Multicast Router.
2. Select Show Static Multicast Router from the Action list.
3. Select the VLAN for which to display this information.

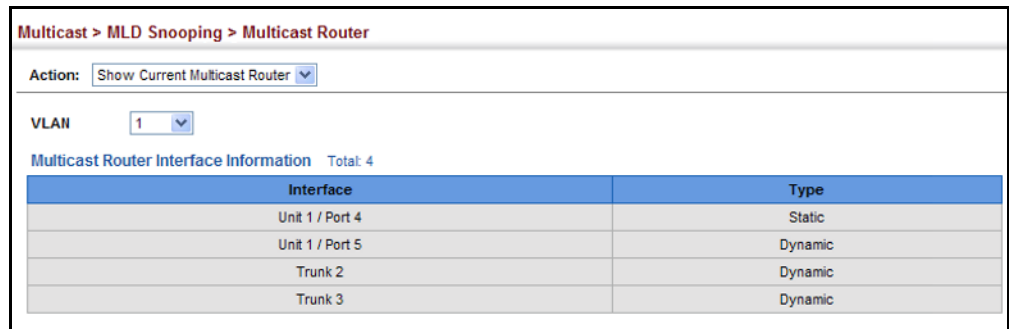
Figure 332: Showing Static Interfaces Attached an IPv6 Multicast Router



To show all the interfaces attached to a multicast router:

1. Click Multicast, MLD Snooping, Multicast Router.
2. Select Current Multicast Router from the Action list.
3. Select the VLAN for which to display this information. Ports in the selected VLAN which are attached to a neighboring multicast router/switch are displayed.

Figure 333: Showing Current Interfaces Attached an IPv6 Multicast Router



Assigning Interfaces to IPv6 Multicast Services

Use the Multicast > MLD Snooping > MLD Member (Add Static Member) page to statically assign an IPv6 multicast service to an interface.

Multicast filtering can be dynamically configured using MLD snooping and query messages (see [“Configuring MLD Snooping and Query Parameters” on page 494](#)). 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.

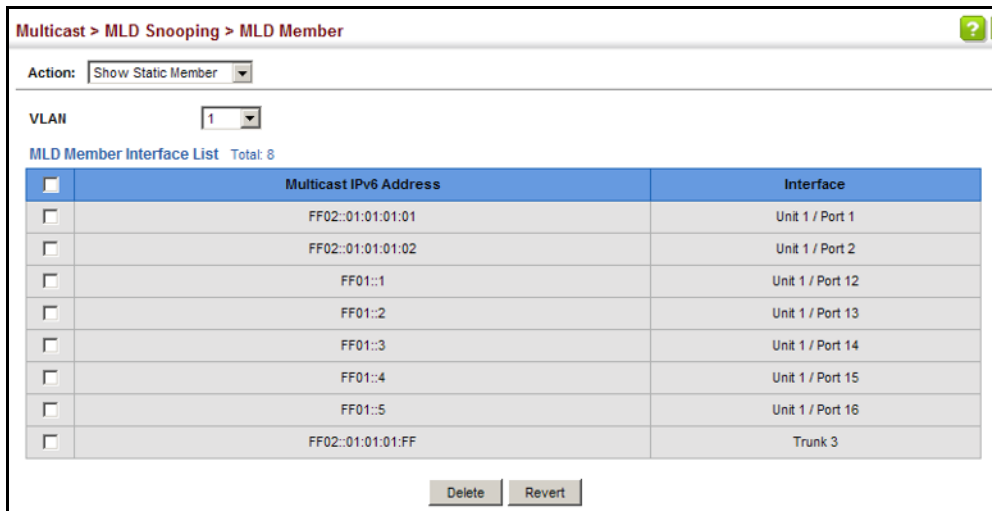
Figure 334: Assigning an Interface to an IPv6 Multicast Service

The screenshot shows a web interface titled "Multicast > MLD Snooping > MLD Member". The "Action" dropdown menu is set to "Add Static Member". The "VLAN" dropdown menu is set to "1". The "Multicast IPv6 Address" text input field contains "FF00:0:0:0:0:10C". The "Interface" section has two radio buttons: "Port" (selected) and "Trunk". The "Port" radio button is selected, and its dropdown menu is set to "1". The "Trunk" radio button is unselected, and its dropdown menu is empty. At the bottom right, there are "Apply" and "Revert" buttons.

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.

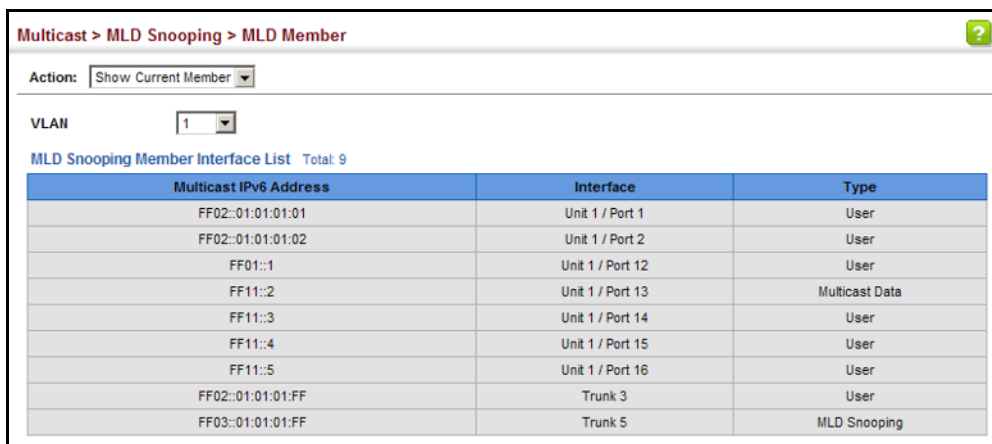
Figure 335: Showing Static Interfaces Assigned to an IPv6 Multicast Service



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 336: Showing Current Interfaces Assigned to an IPv6 Multicast Service



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 only 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 337: Showing IPv6 Multicast Services and Corresponding Sources

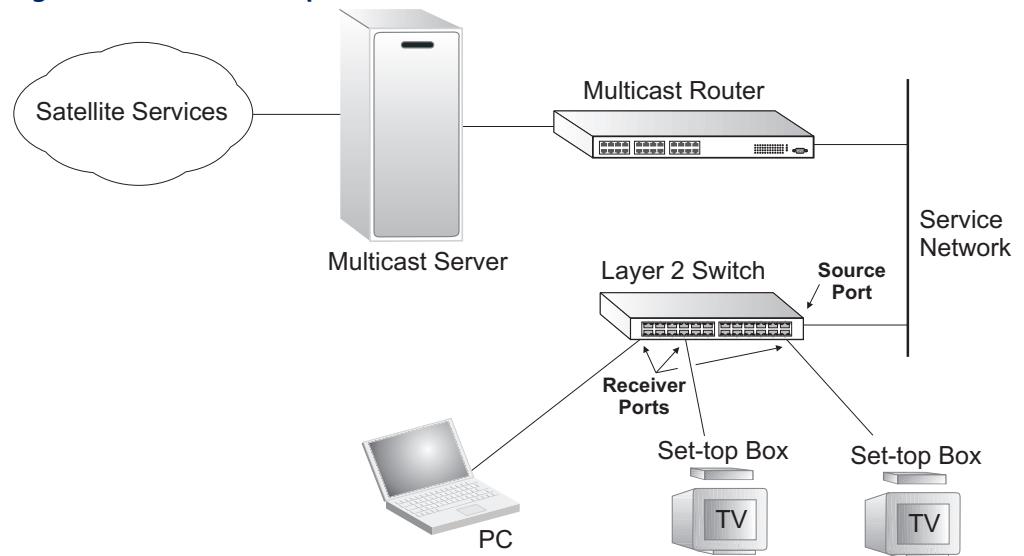


Multicast VLAN Registration

Multicast VLAN Registration (MVR) is a protocol that controls access to a single network-wide VLAN most commonly used for transmitting multicast traffic (such as television channels or video-on-demand) across a service provider's network. Any multicast traffic entering an MVR VLAN is sent to all attached subscribers. This protocol can significantly reduce to processing overhead required to dynamically monitor and establish the distribution tree for a normal multicast VLAN. This makes it possible to support common multicast services over a wide part of the network without having to use any multicast routing protocol.

MVR maintains the user isolation and data security provided by VLAN segregation by passing only multicast traffic into other VLANs to which the subscribers belong. Even though common multicast streams are passed onto different VLAN groups from the MVR VLAN, users in different IEEE 802.1Q or private VLANs cannot exchange any information (except through upper-level routing services).

Figure 338: MVR Concept



Command Usage

◆ General Configuration Guidelines for MVR:

1. Enable MVR globally on the switch, select the MVR VLAN, and add the multicast groups that will stream traffic to attached hosts (see [“Configuring MVR Domain Settings”](#) on page 506).
 2. Set the interfaces that will join the MVR as source ports or receiver ports (see [“Configuring MVR Interface Status”](#) on page 510).
 3. 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 512).
- ◆ 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. 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 512).

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.
 - **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 339: Configuring Global Settings for MVR

The screenshot shows a web interface for configuring Multicast VLAN Registration (MVR) settings. The page title is "Multicast > MVR". Below the title, there is a "Step:" dropdown menu set to "1. Configure Global". The configuration area contains the following settings:

- 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 set to "Always Forward".

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

Configuring MVR Domain Settings Use the Multicast > MVR (Configure Domain) page to enable MVR globally on the switch, 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 140](#)), 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 510](#)).
- ◆ **MVR Current Learned Groups** – The number of MVR groups currently assigned to this domain.
- ◆ **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 global settings for MVR:

1. Click Multicast, MVR.
2. Select Configure General from the Step list.
3. Enable MVR globally on the switch, select the MVR VLAN, and add the multicast groups that will stream traffic to participating hosts.
4. Click Apply.

Figure 340: Configuring Domain 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 '2. Configure Domain'. Below this, several configuration fields are listed:

- Domain ID:** A dropdown menu with the value '1' selected.
- MVR Status:** A checkbox labeled 'Enabled' which is checked.
- MVR VLAN:** A dropdown menu with the value '1' selected.
- MVR Running Status:** A text field with the value 'Inactive'.
- MVR Current Learned Groups:** A text field with the value '0'.
- Upstream Source IP:** A text field with the value '192.100.0.12'.

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

Configuring MVR Group Address Profiles

Use the Multicast > MVR (Configure Profile and Associate Profile) pages to assign the multicast group address for required services to one or more MVR domains.

Command Usage

- ◆ Use the Configure Profile page to statically configure all multicast group addresses that will join the MVR VLAN. Any multicast data associated with an MVR group is sent from all source ports to all receiver ports that have registered to receive data from that multicast group.
- ◆ The IP address range from 224.0.0.0 to 239.255.255.255 is used for multicast streams. MVR group addresses cannot fall within the reserved IP multicast address range of 224.0.0.x.
- ◆ IGMP snooping and MVR share a maximum number of 1024 groups. Any multicast streams received in excess of this limitation will be flooded to all ports in the associated domain.

Parameters

These parameters are displayed:

Configure Profile

- ◆ **Profile Name** – The name of a profile containing one or more MVR group addresses. (Range: 1-21 characters)
- ◆ **Start IP Address** – Starting IP address for an MVR multicast group. (Range: 224.0.1.0 - 239.255.255.255)
- ◆ **End IP Address** – Ending IP address for an MVR multicast group. (Range: 224.0.1.0 - 239.255.255.255)

Associate Profile

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)

- ◆ **Profile Name** – The name of a profile to be assigned to this domain.
(Range: 1-21 characters)

Web Interface

To configure an MVR group address profile:

1. Click Multicast, MVR.
2. Select Configure Profile from the Step list.
3. Select Add from the Action list.
4. Enter the name of a group profile to be assigned to one or more domains, and specify a multicast group that will stream traffic to participating hosts.
5. Click Apply.

Figure 341: Configuring an MVR Group Address Profile

Multicast > MVR

Step: 2. Configure Profile Action: Add

Profile Name: sales

Start IP Address: 234.5.6.8

End IP Address: 234.5.6.10

Apply Revert

To show the configured MVR group address profiles:

1. Click Multicast, MVR.
2. Select Configure Profile from the Step list.
3. Select Show from the Action list.

Figure 342: Displaying MVR Group Address Profiles

Multicast > MVR

Step: 2. Configure Profile Action: Show

MVR Profile List Total: 2

<input type="checkbox"/>	Profile Name	Start IP Address	End IP Address
<input type="checkbox"/>	sales	234.5.6.8	234.5.6.10
<input type="checkbox"/>	rd	234.5.6.7	234.5.6.7

Delete Revert

To assign an MVR group address profile to a domain:

1. Click Multicast, MVR.
2. Select Associate Profile from the Step list.
3. Select Add from the Action list.
4. Select a domain from the scroll-down list, and enter the name of a group profile.
5. Click Apply.

Figure 343: Assigning an MVR Group Address Profile to a Domain

Multicast > MVR

Step: 3. Associate Profile Action: Add

Domain ID: 1

Profile Name: sales

Apply Revert

To show the MVR group address profiles assigned to a domain:

1. Click Multicast, MVR.
2. Select Associate Profile from the Step list.
3. Select Show from the Action list.

Figure 344: Showing the MVR Group Address Profiles Assigned to a Domain

Multicast > MVR

Step: 3. Associate Profile Action: Show

Domain ID: 1

Domain Associated Profile List Total: 2

<input type="checkbox"/>	Profile Name	Start IP Address	End IP Address
<input type="checkbox"/>	rd	234.5.6.7	234.5.6.7
<input type="checkbox"/>	sales	234.5.6.8	234.5.6.10

Delete Revert

Configuring MVR Interface Status Use the Multicast > MVR (Configure Interface) page to configure each interface that participates in the MVR protocol as a source port or receiver port. If you are sure that only one subscriber attached to an interface is receiving multicast services, you can enable the immediate leave function.

Command Usage

- ◆ A port configured as an MVR receiver or source port can join or leave multicast groups configured under MVR. However, note that these ports can also use IGMP snooping to join or leave any other multicast groups using the standard rules for multicast filtering.
- ◆ Receiver ports can belong to different VLANs, but should not be configured as a member of the MVR VLAN. MVR allows a receiver port to dynamically join or leave multicast groups sourced through the MVR VLAN. Multicast groups can also be statically assigned to a receiver port (see [“Assigning Static MVR Multicast Groups to Interfaces” on page 512](#)).

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 140](#)).

- ◆ 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 512](#)).

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** – Port 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 140).
- **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 512).
- **Non-MVR** – An interface that does not participate in the MVR VLAN. (This is the default type.)
- ◆ **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.)

Web Interface

To configure interface settings for MVR:

1. Click Multicast, MVR.
2. Select Configure Interface from the Step list.
3. Select Configure Port or Configure Trunk from the Action list.
4. Select an MVR domain.
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 345: Configuring Interface Settings for MVR

The screenshot shows the 'Multicast > MVR' configuration page. At the top, there are dropdown menus for 'Step: 5. Configure Interface' and 'Action: Configure Port'. Below that is a 'Domain ID' dropdown set to '1'. A 'Port List' section shows 'Total: 28' and three numbered tabs (1, 2, 3). The main table has five columns: Port, Type, Forwarding Status, MVR Status, and Immediate Leave. The rows are as follows:

Port	Type	Forwarding Status	MVR Status	Immediate Leave
1	Source	Forwarding	Inactive	<input type="checkbox"/> Enabled
2	Receiver	Discarding	Inactive	<input checked="" type="checkbox"/> Enabled
3	Non-MVR	Discarding	Inactive	<input type="checkbox"/> Enabled
4	Non-MVR	Discarding	Inactive	<input type="checkbox"/> Enabled
5	Non-MVR	Discarding	Inactive	<input type="checkbox"/> Enabled

Assigning Static MVR Multicast Groups to Interfaces

Use the Multicast > MVR (Configure Static Group Member) page to statically bind multicast groups to a port which will receive long-term multicast streams associated with a stable set of hosts.

Command Usage

- ◆ Multicast groups can be statically assigned to a receiver port using this configuration page.
- ◆ The IP address range from 224.0.0.0 to 239.255.255.255 is used for multicast streams. MVR group addresses cannot fall within the reserved IP multicast address range of 224.0.0.x.
- ◆ Only IGMP version 2 or 3 hosts can issue multicast join or leave messages. If MVR must be configured for an IGMP version 1 host, the multicast groups must be statically assigned.
- ◆ The MVR VLAN cannot be specified as the receiver VLAN for static bindings.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **Interface** – Port or trunk identifier.
- ◆ **VLAN** – VLAN identifier. (Range: 1-4094)
- ◆ **Group IP Address** – Defines a multicast service sent to the selected port. Multicast groups must be assigned from the MVR group range configured on the Configure General page.

Web Interface

To assign a static MVR group to a port:

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 port member to receive the multicast stream, and then enter the multicast group address.
6. Click Apply.

Figure 346: Assigning Static MVR Groups to a Port

The screenshot shows the 'Multicast > MVR' configuration page. At the top, there is a breadcrumb 'Multicast > MVR'. Below it, the 'Step' is set to '5. Configure Static Group Member' and the 'Action' is 'Add'. The configuration fields are as follows:

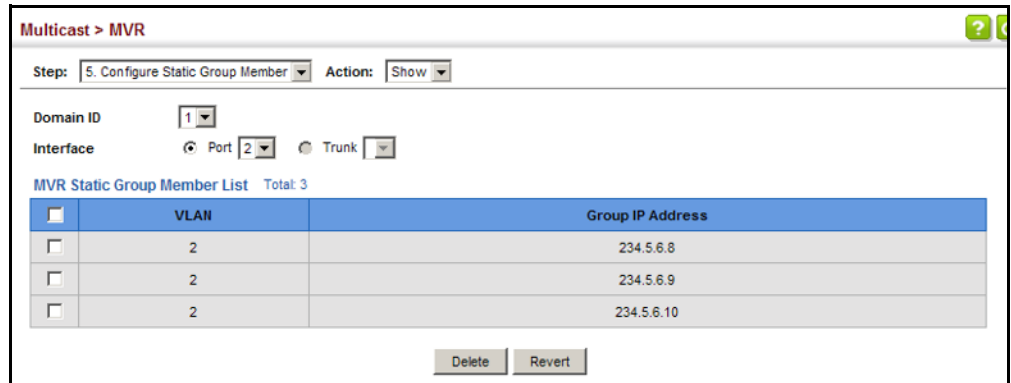
- Domain ID:** A dropdown menu with '1' selected.
- Interface:** Radio buttons for 'Port' (selected) and 'Trunk'. The 'Port' dropdown is set to '2'.
- VLAN:** A dropdown menu with '2' selected.
- Group IP Address:** A text input field containing '234.5.6.8'.

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

To show the static MVR groups assigned to a port:

1. Click Multicast, MVR.
2. Select Configure Static Group Member from the Step list.
3. Select Show from the Action list.
4. Select the port for which to display this information.

Figure 347: Showing the Static MVR Groups Assigned to a Port



Displaying MVR Receiver Groups Use the Multicast > MVR (Show Member) page to show the multicast groups either statically or dynamically assigned to the MVR receiver groups on each interface.

Parameters

These parameters are displayed:

- ◆ **Domain ID** – An independent multicast domain. (Range: 1-5)
- ◆ **Group IP Address** – Multicast groups assigned to the MVR VLAN.
- ◆ **VLAN** – Indicates the MVR VLAN receiving the multicast service. 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.

Web Interface

To show the interfaces associated with multicast groups assigned to the MVR VLAN:

1. Click Multicast, MVR.
2. Select Show Member from the Step list.
3. Select an MVR domain.

Figure 348: Displaying MVR Receiver Groups

Multicast > MVR

Step: 6. Show Member

Domain ID: 1

MVR Member List Total: 8

Group IP Address	VLAN	Port	Up Time	Expire	Count
224.1.1.1	2		00:00:30		2 (Port)
	2	Unit 1 / Port 1 (Source)			
	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: 1-12/28)
- ◆ **Trunk** – Trunk identifier. (Range: 1-8/12)

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 349: 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 is a 'Domain ID' dropdown menu set to '1'. The main content area 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 of the statistics section, there is a 'Clear All' button with a tooltip that reads: 'Click this button to clear all MVR statistics of the domain.' To the right of this section 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 350: Displaying MVR Statistics – VLAN

The screenshot shows the configuration page for Multicast > MVR. At the top, there are two dropdown menus: 'Step: 7. Show Statistics' and 'Action: Show VLAN Statistics'. Below these are two more dropdown menus for 'Domain ID' and 'VLAN', both set to '1'. The main content area is divided into two sections: 'Input Statistics' and 'Output Statistics'. Each section contains a table of statistics with values of 0. A 'Refresh' button is located at the bottom right of the statistics area.

Input Statistics			
Report	0	Drop	0
Leave	0	Join Success	0
G Query	0	Group	0
G(-S)-S Query	0		

Output Statistics	
Report	0
Leave	0
G Query	0
G(-S)-S Query	0

To display 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 351: Displaying MVR Statistics – Port

Multicast > MVR

Step: 7. Show Statistics Action: Show Port Statistics

Domain ID: 1
Port: 1

Input Statistics

Report	0	Drop	0
Leave	0	Join Success	0
G Query	0	Group	0
G(-S)-S Query	0		

Output Statistics

Report	0
Leave	0
G Query	0
G(-S)-S Query	0

Refresh

Section III

Appendices

This section provides additional information and includes these items:

- ◆ [“Software Specifications” on page 523](#)
- ◆ [“Troubleshooting” on page 527](#)
- ◆ [“License Information” on page 529](#)



Software Specifications

Software Features

Management Authentication Local, RADIUS, TACACS+, Port Authentication (802.1X), HTTPS, SSH, Port Security, IP Filter, DHCP Snooping

Client Access Control Access Control Lists (512 rules), Port Authentication (802.1X), MAC Authentication, Port Security, DHCP Snooping, IP Source Guard

Port Configuration 1000BASE-T: 10/100 Mbps at half/full duplex, 1000 Mbps at full duplex
1000BASE-SX/LX/LH - 1000 Mbps at 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 50 sessions, one or more source ports to one destination port

Rate Limits Input/Output Limits
Range configured per port

Port Trunking Static trunks (Cisco EtherChannel compliant)
Dynamic trunks (Link Aggregation Control Protocol)

Spanning Tree Algorithm Spanning Tree Protocol (STP, IEEE 802.1D-2004)
Rapid Spanning Tree Protocol (RSTP, IEEE 802.1D-2004)
Multiple Spanning Tree Protocol (MSTP, IEEE 802.1D-2004)

VLAN Support Up to 256 groups; port-based, protocol-based, tagged (802.1Q), private VLANs, voice VLANs, IP subnet, MAC-based, GVRP for automatic VLAN learning

Class of Service Supports four levels of priority
Strict, Shaped Deficit Weighted Round Robin, or strict-WRR queuing
Layer 3/4 priority mapping: IP DSCP

Quality of Service DiffServ (IPv4/v6) supports class maps, policy maps, and service policies

Multicast Filtering IGMP Snooping (Layer 2)
Multicast VLAN Registration

Additional Features BOOTP Client
DHCP Client
DNS Client, Proxy
LLDP (Link Layer Discover Protocol)
RMON (Remote Monitoring, groups 1,2,3,9)
SMTP Email Alerts
SNMP (Simple Network Management Protocol)
SNTP (Simple Network Time Protocol)

Management Features

In-Band Management Telnet, web-based HTTP or HTTPS, SNMP manager, or Secure Shell

Out-of-Band Management RS-232 DB-9 console port

Software Loading HTTP, FTP or TFTP in-band, or XModem out-of-band

SNMP Management access via MIB database
Trap management to specified hosts

RMON Groups 1, 2, 3, 9 (Statistics, History, Alarm, Event)

Standards

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
DHCP Client (RFC 2131)
DHCPv6 Client (RFC 3315)
HTTPS
ICMP (RFC 792)
IGMP (RFC 1112)
IGMPv2 (RFC 2236)
IGMPv3 (RFC 3376) - partial support
IPv4 IGMP (RFC 3228)
RADIUS+ (RFC 2618)
RMON (RFC 2819 groups 1,2,3,9)
SNMP (RFC 1157)
SNMPv2c (RFC 1901, 2571)
SNMPv3 (RFC DRAFT 2273, 2576, 3410, 3411, 3413, 3414, 3415)
SNTP (RFC 2030)
SSH (Version 2.0)
TELNET (RFC 854, 855, 856)
TFTP (RFC 1350)

Management Information Bases

Bridge MIB (RFC 1493)
DHCP Option for Civic Addresses Configuration Information (RFC 4776)
Differentiated Services MIB (RFC 3289)
DNS Resolver MIB (RFC 1612)
Entity MIB (RFC 2737)
Ether-like MIB (RFC 2665)
Extended Bridge MIB (RFC 2674)

Extensible SNMP Agents MIB (RFC 2742)
Forwarding Table MIB (RFC 2096)
IGMP MIB (RFC 2933)
Interface Group MIB (RFC 2233)
Interfaces Evolution MIB (RFC 2863)
IP Multicasting related MIBs
IPV6-MIB (RFC 2065)
IPV6-ICMP-MIB (RFC 2066)
IPV6-TCP-MIB (RFC 2052)
IPV6-UDP-MIB (RFC2054)
Link Aggregation MIB (IEEE 802.3ad)
MAU MIB (RFC 3636)
MIB II (RFC 1213)
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)
Quality of Service MIB
RADIUS Authentication Client MIB (RFC 2621)
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 36: Troubleshooting Chart

Symptom	Action
Cannot connect using Telnet, web browser, or SNMP software	<ul style="list-style-type: none">◆ Be sure the switch is powered up.◆ Check network cabling between the management station and the switch.◆ Check that you have a valid network connection to the switch and that the port you are using has not been disabled.◆ Be sure you have configured the VLAN interface through which the management station is connected with a valid IP address, subnet mask and default gateway.◆ Be sure the management station has an IP address in the same subnet as the switch's IP interface to which it is connected.◆ If you are trying to connect to the switch via the IP address for a tagged VLAN group, your management station, and the ports connecting intermediate switches in the network, must be configured with the appropriate tag.◆ If you cannot connect using Telnet, you may have exceeded the maximum number of concurrent Telnet/SSH sessions permitted. Try connecting again at a later time.
Cannot connect using Secure Shell	<ul style="list-style-type: none">◆ If you cannot connect using SSH, you may have exceeded the maximum number of concurrent Telnet/SSH sessions permitted. Try connecting again at a later time.◆ Be sure the control parameters for the SSH server are properly configured on the switch, and that the SSH client software is properly configured on the management station.◆ Be sure you have generated both an RSA and DSA public key on the switch, exported this key to the SSH client, and enabled SSH service.◆ Be sure you have set up an account on the switch for each SSH user, including user name, authentication level, and password.◆ Be sure you have imported the client's public key to the switch (if public key authentication is used).
Cannot access the on-board configuration program via a serial port connection	<ul style="list-style-type: none">◆ Be sure you have set the terminal emulator program to VT100 compatible, 8 data bits, 1 stop bit, no parity, and the baud rate set to 115200 bps).◆ Check that the null-modem serial cable conforms to the pin-out connections provided in the Installation Guide.
Forgot or lost the password	<ul style="list-style-type: none">◆ Contact your local distributor.

Using System Logs

If a fault does occur, refer to the Installation Guide to ensure that the problem you encountered is actually caused by the switch. If the problem appears to be caused by the switch, follow these steps:

1. Enable logging.
2. Set the error messages reported to include all categories.
3. Enable SNMP.
4. Enable SNMP traps.
5. Designate the SNMP host that is to receive the error messages.
6. Repeat the sequence of commands or other actions that lead up to the error.
7. Make a list of the commands or circumstances that led to the fault. Also make a list of any error messages displayed.
8. Set up your terminal emulation software so that it can capture all console output to a file. Then enter the “show tech-support” command to record all system settings in this file.
9. Contact your distributor’s service engineer, and send a detailed description of the problem, along with the file used to record your system settings.

For example:

```
Console(config)#logging on
Console(config)#logging history flash 7
Console(config)#snmp-server host 192.168.1.23
:
```




License 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. BOOTP is used to provide bootup information for network devices, including IP address information, the address of the TFTP server that contains the devices system files, and the name of the boot file.

CoS Class of Service is supported by prioritizing packets based on the required level of service, and then placing them in the appropriate output queue. Data is transmitted from the queues using weighted round-robin service to enforce priority service and prevent blockage of lower-level queues. Priority may be set according to the port default, the packet's priority bit (in the VLAN tag), TCP/UDP port number, IP Precedence bit, or DSCP priority bit.

DHCP Dynamic Host Control Protocol. Provides a framework for passing configuration information to hosts on a TCP/IP network. DHCP is based on the Bootstrap Protocol (BOOTP), adding the capability of automatic allocation of reusable network addresses and additional configuration options.

DHCP 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 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 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 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.

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.

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.

MRD Multicast Router Discovery is a protocol used by IGMP snooping and multicast routing devices to discover which interfaces are attached to multicast routers. This process allows IGMP-enabled devices to determine where to send multicast source and group membership messages.

Multicast Switching A process whereby the switch filters incoming multicast frames for services for which no attached host has registered, or forwards them to all ports contained within the designated multicast VLAN group.

MVR Multicast VLAN Registration is a method of using a single network-wide multicast VLAN to transmit common services, such as television channels or video-on-demand, across a service-provider's network. MVR simplifies the configuration of multicast services by using a common VLAN for distribution, while still preserving security and data isolation for subscribers residing in both the MVR VLAN and other standard 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.

Out-of-Band Management Management of the network from a station not attached to the network.

Port Authentication See *IEEE 802.1X*.

Port Mirroring A method whereby data on a target port is mirrored to a monitor port for troubleshooting with a logic analyzer or RMON probe. This allows data on the target port to be studied unobstructively.

Port Trunk Defines a network link aggregation and trunking method which specifies how to create a single high-speed logical link that combines several lower-speed physical links.

QinQ QinQ tunneling is designed for service providers carrying traffic for multiple customers across their networks. It is used to maintain customer-specific VLAN and Layer 2 protocol configurations even when different customers use the same internal VLAN IDs.

QoS Quality of Service. QoS refers to the capability of a network to provide better service to selected traffic flows using features such as data prioritization, queuing, congestion avoidance and traffic shaping. These features effectively provide preferential treatment to specific flows either by raising the priority of one flow or limiting the priority of another flow.

- RADIUS** Remote Authentication Dial-in User Service. RADIUS is a logon authentication protocol that uses software running on a central server to control access to RADIUS-compliant devices on the network.
- 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.
- XModem** A protocol used to transfer files between devices. Data is grouped in 128-byte blocks and error-corrected.

Index

Numerics

- 802.1Q tunnel 148
 - access 153
 - configuration, guidelines 151
 - configuration, limitations 151
 - description 148
 - ethernet type 152
 - interface configuration 153
 - mode selection 153
 - status, configuring 152
 - TPID 152
 - uplink 153
- 802.1X
 - authenticator, configuring 326
 - global settings 325
 - port authentication 323
 - port authentication accounting 257, 258
 - supplicant, configuring 330

A

- AAA
 - accounting 802.1X port settings 257, 258
 - accounting exec command privileges 258
 - accounting exec settings 257
 - accounting summary 258
 - accounting update 257
 - accounting, configuring 257
 - authorization & accounting 250
 - authorization exec settings 262
 - authorization method 263
 - authorization settings 262
 - authorization summary 264
 - RADIUS group settings 258
 - TACACS+ group settings 258
- acceptable frame type 141
- ACL 290
 - ARP 295, 306
 - binding to a port 308
 - IPv4 Extended 295, 298
 - IPv4 Standard 295, 297
 - IPv6 Extended 295, 302
 - IPv6 Standard 295, 300
 - MAC 295, 304
 - time range 291
- Address Resolution Protocol *See* ARP

- address table 165
 - address isolation 167
 - aging time 170
 - aging time, displaying 170
 - aging time, setting 170
 - movable static 175
 - sticky dynamic 175
- ARP
 - ACL 306, 315
 - configuration 429
 - description 429
- ARP inspection 312
 - ACL filter 315
 - additional validation criteria 314
 - ARP ACL 315
 - enabling globally 314
 - enabling per VLAN 315
 - trusted ports 316
- ATC 206
 - control response 209
 - functional limitations 207
 - limiting traffic rates 206
 - shutting down a port 207
 - thresholds 209
 - timers 208
 - usage 206
- authentication
 - MAC address authentication 270
 - MAC, configuring ports 273
 - network access 270
 - public key 284
 - web 267
 - web authentication port information, displaying 268
 - web authentication, configuring ports 268
 - web authentication, re-authenticating address 269
 - web authentication, re-authenticating ports 268
 - web, configuring 267
- Automatic Traffic Control *See* ATC

B

- BOOTP 431
- BPDU 180
 - filter 192
 - flooding when STA disabled on VLAN 185
 - flooding when STA globally disabled 185

Index

- guard 192
- ignoring superior BPDUs 191
- selecting protocol based on message format 193
- shut down port on receipt 192

bridge extension capabilities, displaying 69

broadcast storm, threshold 204, 205

C

cable diagnostics 110

canonical format indicator 224

class map

- DiffServ 228

Class of Service *See* CoS

clustering switches, management access 421

committed burst size, QoS policy 236, 237

committed information rate, QoS policy 236, 237

community string 398

configuration files, restoring defaults 71

configuration settings

- restoring 73, 74
- saving 73

CoS 213

- configuring 213
- default mapping to internal values 224
- enabling 220
- layer 3/4 priorities 220
- priorities, mapping to internal values 224
- queue mapping 217
- queue mode 214
- queue weights, assigning 215

CoS/CFI to PHB/drop precedence 224

CPU

- status 90
- utilization, showing 90

D

default IPv4 gateway, configuration 432

default IPv6 gateway, configuration 434

default priority, ingress port 213

default settings, system 39

DHCP 431, 460

- class identifier 461
- client 431
- client identifier 460, 461
- option 82 information 461
- relay option 82 461
- relay server 461

DHCP snooping 348

- enabling 348
- global configuration 348
- information option 348
- information option policy 349
- information option, enabling 348

- policy selection 349
- specifying trusted interfaces 350
- verifying MAC addresses 348
- VLAN configuration 350

DHCPv4 snooping

- option 82 information, circuit-id 351

Differentiated Code Point Service *See* DSCP

Differentiated Services *See* DiffServ

DiffServ 227

- binding policy to interface 241
- class map 228
- classifying QoS traffic 228
- color aware, srTCM 237
- color aware, trTCM 237
- color blind, srTCM 237
- color blind, trTCM 237
- committed burst size 236, 237, 238
- committed information rate 236, 237, 238
- configuring 227
- conforming traffic, configuring response 236
- excess burst size 237
- metering, configuring 232, 233
- peak burst size 238
- peak information rate 238
- policy map 232
- policy map, description 228, 235
- QoS policy 232
- service policy 241
- setting CoS for matching packets 235
- setting IP DSCP for matching packets 236
- setting PHB for matching packets 235
- single-rate, three-color meter 232, 236
- srTCM metering 232, 236
- traffic between CIR and BE, configuring response 236
- traffic between CIR and PIR, configuring response 237
- trTCM metering 237
- two-rate, three-color meter 233
- violating traffic, configuring response 238

DNS

- default domain name 453
- displaying the cache 458
- domain name list 453
- enabling lookup 453
- multicast name service 459
- name server list 453
- static entries, IPv4 457

Domain Name Service *See* DNS

DoS protection 335

downloading software 71

- automatically 75
- using FTP or TFTP 75

drop precedence

- CoS priority mapping 224, 235, 236
- DSCP ingress map 222

DSA encryption 286, 288
 DSCP 220
 enabling 220
 mapping to internal values 221
 DSCP ingress map, drop precedence 222
 DSCP to PHB/drop precedence 222
 dynamic addresses
 clearing 172
 displaying 171
 dynamic QoS assignment 271, 274
 dynamic VLAN assignment 270, 274

E

edge port, STA 192, 195
 encryption
 DSA 286, 288
 RSA 286, 288
 engine ID 388, 389
 event logging 355
 excess burst size, QoS policy 236
 exec command privileges, accounting 258
 exec settings
 accounting 257
 authorization 262

F

firmware
 displaying version 67
 upgrading 71
 upgrading automatically 75
 upgrading with FTP or TFTP 75
 version, displaying 67

G

GARP VLAN Registration Protocol *See* GVRP
 gateway, IPv4 default 432
 gateway, IPv6 default 434
 general security measures 249
 GNU license 529
 GVRP
 enabling 145
 global setting 145
 interface configuration 145

H

hardware version, displaying 67
 HTTPS 279, 281
 configuring 279
 replacing SSL certificate 281
 secure-site certificate 281
 UDP port, configuring 280

HTTPS, secure server 279

I

IEEE 802.1D 179
 IEEE 802.1s 179
 IEEE 802.1w 179
 IEEE 802.1X 323
 IGMP
 filter profiles, binding to interface 492
 filter profiles, configuration 490
 filter, interface configuration 492
 filter, parameters 490
 filtering & throttling 489
 filtering & throttling, enabling 489
 filtering & throttling, interface configuration 492
 filtering & throttling, status 489
 filtering, configuring profile 490
 filtering, creating profile 490
 filtering, group range 490
 groups, displaying 477
 Layer 2 468
 query 468, 470
 query, enabling 473
 services, displaying 484
 snooping 468
 snooping & query, parameters 470
 snooping, configuring 470
 snooping, enabling 470
 snooping, immediate leave 479
 IGMP snooping
 configuring 478
 enabling per interface 478, 479
 forwarding entries 484
 immediate leave, status 479
 interface attached to multicast router 475
 last leave 469
 last member query count 481
 last member query interval 481
 proxy query address 482
 proxy reporting 470, 480
 querier timeout 472
 querier, enabling 473
 query interval 481
 query response interval 481
 query suppression 469
 router port expire time 472
 static host interface 469
 static multicast routing 473
 static port assignment 476
 static router interface 469
 static router port, configuring 473
 statistics, displaying 485
 TCN flood 471
 unregistered data flooding 472

Index

- version exclusive 472
- version for interface, setting 481
- version, setting 472
- with proxy reporting 469
- immediate leave, IGMP snooping 479
- immediate leave, MLD snooping 496
- importing user public keys 288
- ingress filtering 141
- IP address
 - Auto IP 431
 - BOOTP/DHCP 431
 - setting 427
- IP filter, for management access 319
- IP source guard
 - configuring static entries 343
 - setting filter criteria 341
 - setting maximum bindings 342
- IPv4 address
 - BOOTP/DHCP 431
 - setting 431
- IPv6
 - displaying neighbors 444
 - duplicate address detection 437, 444
 - enabling 436
 - MTU 436
 - statistics 445
- IPv6 address
 - dynamic configuration (global unicast) 441
 - dynamic configuration (link-local) 436
 - EUI format 441
 - EUI-64 setting 441
 - explicit configuration 436
 - global unicast 441
 - link-local 442
 - manual configuration (global unicast) 441
 - manual configuration (link-local) 442
 - setting 434

J

jumbo frame 68

K

key

- private 282
- public 282
- user public, importing 288

key pair

- host 282
- host, generating 286

L

LACP

- configuration 118
- group attributes, configuring 121
- group members, configuring 119
- local parameters 124
- partner parameters 126
- protocol message statistics 123
- protocol parameters 118

last member query count, IGMP snooping 481

last member query interval, IGMP snooping 481

license information, GNU 529

Link Layer Discovery Protocol - Media Endpoint Discovery
See LLDP-MED

Link Layer Discovery Protocol See LLDP

link type, STA 191, 195

LLDP 359

- device statistics details, displaying 381
- device statistics, displaying 379
- display device information 367, 371
- displaying remote information 371
- interface attributes, configuring 361
- local device information, displaying 367
- message attributes 361
- message statistics 379
- remote information, displaying 378, 379
- remote port information, displaying 371
- timing attributes, configuring 360
- TLV 359, 362
- TLV, 802.1 363
- TLV, 802.3 363
- TLV, basic 362
- TLV, management address 362
- TLV, port description 362
- TLV, system capabilities 363
- TLV, system description 363
- TLV, system name 363

LLDP-MED 359

- notification, status 362
- TLV 364
- TLV, extended PoE 364
- TLV, inventory 364
- TLV, location 364
- TLV, MED capabilities 364
- TLV, network policy 364
- TLV, PoE 364

local engine ID 388

logging

- messages, displaying 357
- syslog traps 358
- to syslog servers 358

log-in, web interface 44

logon authentication 265

encryption keys 254

- RADIUS client 253
- RADIUS server 253
- sequence 251
- settings 252
- TACACS+ client 252
- TACACS+ server 252
- logon authentication, settings 254
- loopback detection, STA 182

M

- MAC address authentication 270
 - ports, configuring 273
 - reauthentication 272
- MAC address, mirroring 173
- main menu, web interface 47
- management access, filtering per address 319
- management access, IP filter 319
- Management Information Bases (MIBs) 525
- matching class settings, classifying QoS traffic 229
- memory
 - status 91
 - utilization, showing 91
- mirror port
 - configuring 99
 - configuring local traffic 99
 - configuring remote traffic 101
- MLD snooping 494
 - configuring 494
 - enabling 494
 - groups, displaying 500, 501
 - immediate leave 496
 - immediate leave, status 496
 - interface attached to multicast router 496, 498
 - multicast static router port 496
 - querier 494
 - querier, enabling 494
 - query interval 494
 - query, maximum response time 495
 - robustness value 494
 - static port assignment 498
 - static router port 496
 - unknown multicast, handling 495
 - version 495
- movable static address 175
- MSTP 179, 196
 - global settings, configuring 183, 196
 - global settings, displaying 189
 - interface settings, configuring 190, 200
 - interface settings, displaying 202
 - max hop count 186
 - path cost 201
 - region name 186
 - region revision 186
- MTU for IPv6 436

- multicast filtering 467
 - enabling IGMP snooping 470, 479
 - enabling IGMP snooping per interface 478
 - enabling MLD snooping 494
 - router configuration 473
- multicast groups 477, 484, 500
 - displaying 477, 484, 500
 - static 476, 477, 498, 500
- multicast router discovery 478
- multicast router port, displaying 475, 498
- multicast services
 - configuring 476, 498
 - displaying 477, 500
- multicast static router port 473
 - configuring 473
 - configuring for MLD snooping 496
- multicast storm, threshold 204, 205
- Multicast VLAN Registration *See* MVR
- multicast, filtering and throttling 489
- MVR
 - assigning static multicast groups 512
 - configuring 506
 - description 502
 - interface status, configuring 510
 - interface status, displaying 511
 - IP for control packets sent upstream 506
 - proxy switching 504
 - robust value for proxy switching 504
 - setting interface type 510
 - setting multicast domain 506
 - setting multicast groups 506, 507
 - specifying a domain 506
 - specifying a VLAN 506
 - static binding 512
 - static binding, group to port 512
 - statistics, displaying 515
 - using immediate leave 511

N

- network access
 - authentication 270
 - dynamic QoS assignment 274
 - dynamic VLAN assignment 274
 - guest VLAN 274
 - MAC address filter 274
 - port configuration 273
 - reauthentication 272
 - secure MAC information 278
- NTP
 - authentication keys, specifying 84
 - client, enabling 81
 - setting the system clock 83
 - specifying servers 83

Index

P

- passwords
 - administrator setting 265
- path cost 195
 - method 185
 - STA 191, 195
- peak burst size, QoS policy 237
- peak information rate, QoS policy 237
- per-hop behavior, DSCP ingress map 222
- policing traffic, QoS policy 232, 236
- policy map
 - description 235
 - DiffServ 232
- port authentication 323
- port power
 - displaying status 384
 - inline 383
 - inline status 384
 - maximum allocation 383
 - priority 384
 - showing main power 384
- port priority
 - configuring 213
 - default ingress 213
 - STA 190
- port security, configuring 321
- ports
 - autonegotiation 96
 - broadcast storm threshold 204, 205
 - capabilities 96
 - configuring 95
 - duplex mode 97
 - flow control 97
 - mirroring 99
 - mirroring local traffic 99
 - mirroring remote traffic 101
 - multicast storm threshold 204, 205
 - speed 97
 - statistics 105
 - unknown unicast storm threshold 204, 205
- power budgets
 - port 383
 - port priority 384
- priority, default port ingress 213
- private key 282
- problems, troubleshooting 527
- protocol migration 193
- protocol VLANs 154
 - configuring 155
 - configuring groups 155
 - configuring interfaces 156
 - group configuration 155
 - interface configuration 156
- proxy query address, IGMP snooping 482

- proxy reporting, IGMP snooping 480
- public key 282
- PVID, port native VLAN 141

Q

- QinQ Tunneling *See* 802.1Q tunnel
- QoS 227
 - configuration guidelines 228
 - configuring 227
 - CoS/CFI to PHB/drop precedence 224
 - DSCP to PHB/drop precedence 221
 - dynamic assignment 274
 - matching class settings 229
 - PHB to queue 217
 - selecting DSCP, CoS 220
- QoS policy
 - committed burst size 236, 237
 - excess burst size 236
 - peak burst size 237
 - policing flow 232, 236
 - srTCM 232
 - srTCM police meter 236
 - trTCM 233
 - trTCM police meter 237
- QoS policy, committed information rate 236, 237
- QoS policy, peak information rate 237
- Quality of Service *See* QoS
- query interval, IGMP snooping 481
- query response interval, IGMP snooping 481
- queue mode, setting 214
- queue weight, assigning to CoS 215

R

- RADIUS
 - logon authentication 253
 - settings 253
- rate limit
 - port 203
 - setting 203
- remote engine ID 389
- remote logging 358
- restarting the system 91
 - at scheduled times 91
- RMON 411
 - alarm, displaying settings 414
 - alarm, setting thresholds 412
 - event settings, displaying 416
 - response to alarm setting 414
 - statistics history, collection 416
 - statistics history, displaying 418
 - statistics, collection 419
 - statistics, displaying 420
- root guard 191

- RSA encryption 286, 288
 - RSTP 179
 - global settings, configuring 183
 - global settings, displaying 189
 - interface settings, configuring 190
 - interface settings, displaying 194
- S**
- secure shell 282
 - configuration 282
 - security, general measures 249
 - serial port, configuring 86
 - Simple Network Management Protocol *See* SNMP
 - single rate three color meter *See* srTCM
 - SNMP 385
 - community string 398
 - enabling traps 403
 - filtering IP addresses 319
 - global settings, configuring 387
 - trap manager 403
 - users, configuring 399, 401
 - SNMPv3
 - engine ID 388, 389
 - engine identifier, local 388
 - engine identifier, remote 389
 - groups 393
 - local users, configuring 399
 - remote users, configuring 401
 - user configuration 399, 401
 - views 390
 - SNTP
 - setting the system clock 80
 - specifying servers 82
 - software
 - displaying version 67
 - downloading 71
 - version, displaying 67
 - Spanning Tree Protocol *See* STA
 - specifications, software 523
 - srTCM
 - police meter 236
 - QoS policy 232
 - SSH 282
 - authentication retries 285
 - configuring 282
 - downloading public keys for clients 288
 - generating host key pair 286
 - server, configuring 285
 - timeout 285
 - SSL, replacing certificate 281
 - STA 179
 - BPDU filter 192
 - BPDU flooding 185, 190
 - BPDU shutdown 192
 - detecting loopbacks 182
 - edge port 192, 195
 - forward delay 186
 - global settings, configuring 183
 - global settings, displaying 189
 - hello time 185
 - interface settings, configuring 190
 - interface settings, displaying 194
 - link type 191, 195
 - loopback detection 182
 - maximum age 185
 - MSTP interface settings, configuring 200
 - MSTP path cost 201
 - path cost 191, 195
 - path cost method 185
 - port priority 190
 - port/trunk loopback detection 182
 - protocol migration 193
 - transmission limit 185
 - standards, IEEE 525
 - startup files
 - creating 71
 - displaying 71
 - setting 71
 - static addresses, setting 165
 - statistics, port 105
 - sticky dynamic address 175
 - STP 184
 - summary, accounting 258
 - switch clustering, for management 421
 - switch settings
 - restoring 73
 - saving 73
 - system clock
 - setting 79
 - setting manually 79
 - setting the time zone 85
 - setting with NTP 83
 - setting with SNTP 80
 - system logs 355
 - system software, downloading from server 71
- T**
- TACACS+
 - logon authentication 252
 - settings 254
 - TCN
 - flood 471
 - general query solicitation 471
 - Telnet
 - configuring 88
 - server, enabling 88
 - telnet connection, configuring 88
 - time range, ACL 291

Index

- time zone, setting 85
 - time, setting 79
 - TPID 152
 - traffic segmentation 128
 - assigning ports 128
 - enabling 128
 - sessions, assigning ports 130
 - sessions, creating 129
 - trap manager 403
 - troubleshooting 527
 - trTCM
 - police meter 237
 - QoS policy 233
 - trunk
 - configuration 115
 - LACP 118
 - static 116
 - tunneling unknown VLANs, VLAN trunking 131
 - two rate three color meter *See* trTCM
 - Type Length Value
 - See* LLDP TLV
- ## U
- unknown unicast storm, threshold 204, 205
 - unregistered data flooding, IGMP snooping 472
 - upgrading software 71
 - user account 265
 - user password 265
- ## V
- VLAN trunking 131
 - VLANs 135–162
 - 802.1Q tunnel mode 153
 - acceptable frame type 141
 - adding static members 140
 - configuring port members, VLAN index 143
 - creating 138
 - description 135
 - displaying port members 143
 - displaying port members by interface 144
 - displaying port members by interface range 144
 - displaying port members by VLAN index 143
 - dynamic assignment 274
 - egress mode 141
 - ingress filtering 141
 - interface configuration 140
 - IP subnet-based 158
 - MAC-based 160
 - mirroring 162
 - port members, displaying 143
 - protocol 154
 - protocol, configuring 155
 - protocol, configuring groups 155
 - protocol, configuring interfaces 156
 - protocol, group configuration 155
 - protocol, interface configuration 156
 - PVID 141
 - tunneling unknown groups 131
 - voice 243
 - voice VLANs 243
 - detecting VoIP devices 244
 - enabling for ports 247
 - identifying client devices 245
 - VoIP traffic 243
 - ports, configuring 246
 - telephony OUI, configuring 245
 - voice VLAN, configuring 244
 - VoIP, detecting devices 247
- ## W
- web authentication 267
 - address, re-authenticating 269
 - configuring 267
 - configuring ports 268
 - port information, displaying 268, 269
 - ports, configuring 268
 - ports, re-authenticating 268
 - web interface
 - access requirements 43
 - configuration buttons 45
 - home page 44
 - menu list 47
 - panel display 46

ECS4210-12P	149100000219A
ECS4210-12T	149100000241A
ECS4210-28P	149100000217H
ECS4210-28T	149100000217H

E032014/ST-R03