



INSTRUCTION MANUAL

NXA-ENET8-POE+

GIGABIT POE ETHERNET SWITCH



IMPORTANT SAFETY INSTRUCTIONS

1. READ these instructions.
2. KEEP these instructions.
3. HEED all warnings.
4. FOLLOW all instructions.
5. DO NOT use this apparatus near water.
6. CLEAN ONLY with dry cloth.
7. DO NOT block any ventilation openings. Install in accordance with the manufacturer's instructions.
8. DO NOT install near any heat sources such as radiators, heat registers, stoves, or other apparatus (including amplifiers) that produce heat.
9. DO NOT defeat the safety purpose of the polarized or grounding type plug. A polarized plug has two blades with one wider than the other. A grounding type plug has two blades and a third grounding prong. The wider blade or the third prong are provided for your safety. If the provided plug does not fit into your outlet, consult an electrician for replacement of the obsolete outlet.
10. PROTECT the power cord from being walked on or pinched, particularly at plugs, convenience receptacles, and the point where they exit from the apparatus.
11. ONLY USE attachments/accessories specified by the manufacturer.



12. USE ONLY with a cart, stand, tripod, bracket, or table specified by the manufacturer, or sold with the apparatus. When a cart is used, use caution when moving the cart/apparatus combination to avoid injury from tip-over.
13. UNPLUG this apparatus during lightning storms or when unused for long periods of time.
14. REFER all servicing to qualified service personnel. Servicing is required when the apparatus has been damaged in any way, such as power-supply cord or plug is damaged, liquid has been spilled or objects have fallen into the apparatus, the apparatus has been exposed to rain or moisture, does not operate normally, or has been dropped.
15. DO NOT expose this apparatus to dripping or splashing and ensure that no objects filled with liquids, such as vases, are placed on the apparatus.
16. To completely disconnect this apparatus from the AC Mains, disconnect the power supply cord plug from the AC receptacle.
17. Where the mains plug or an appliance coupler is used as the disconnect device, the disconnect device shall remain readily operable.
18. DO NOT overload wall outlets or extension cords beyond their rated capacity as this can cause electric shock or fire.



The exclamation point, within an equilateral triangle, is intended to alert the user to the presence of important operating and maintenance (servicing) instructions in the literature accompanying the product.



The lightning flash with arrowhead symbol within an equilateral triangle is intended to alert the user to the presence of uninsulated "dangerous voltage" within the product's enclosure that may be of sufficient magnitude to constitute a risk of electrical shock to persons.



ESD Warning: The icon to the left indicates text regarding potential danger associated with the discharge of static electricity from an outside source (such as human hands) into an integrated circuit, often resulting in damage to the circuit.

- WARNING:** To reduce the risk of fire or electrical shock, do not expose this apparatus to rain or moisture.
- WARNING:** No naked flame sources - such as lighted candles - should be placed on the product.
- WARNING:** Equipment shall be connected to a MAINS socket outlet with a protective earthing connection.



WARNING: The bottom of the enclosure is a hot surface. Do not touch!

COPYRIGHT NOTICE

AMX© 2016, all rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of AMX. Copyright protection claimed extends to AMX hardware and software and includes all forms and matters copyrightable material and information now allowed by statutory or judicial law or herein after granted, including without limitation, material generated from the software programs which are displayed on the screen such as icons, screen display looks, etc. Reproduction or disassembly of embodied computer programs or algorithms is expressly prohibited.


LIABILITY NOTICE


No patent liability is assumed with respect to the use of information contained herein. While every precaution has been taken in the preparation of this publication, AMX assumes no responsibility for error or omissions. No liability is assumed for damages resulting from the use of the information contained herein. Further, this publication and features described herein are subject to change without notice.

AMX WARRANTY AND RETURN POLICY

The AMX Warranty and Return Policy and related documents can be viewed/downloaded at www.amx.com.

ESD WARNING

	<p>To avoid ESD (Electrostatic Discharge) damage to sensitive components, make sure you are properly grounded before touching any internal materials.</p> <p>When working with any equipment manufactured with electronic devices, proper ESD grounding procedures must be followed to make sure people, products, and tools are as free of static charges as possible. Grounding straps, conductive smocks, and conductive work mats are specifically designed for this purpose. These items should not be manufactured locally, since they are generally composed of highly resistive conductive materials to safely drain static discharges, without increasing an electrocution risk in the event of an accident.</p> <p>Anyone performing field maintenance on AMX equipment should use an appropriate ESD field service kit complete with at least a dissipative work mat with a ground cord and a UL listed adjustable wrist strap with another ground cord.</p>
---	---

	<p>WARNING: Do Not Open! Risk of Electrical Shock. Voltages in this equipment are hazardous to life. No user-serviceable parts inside. Refer all servicing to qualified service personnel.</p> <p>Place the equipment near a main power supply outlet and make sure that you can easily access the power breaker switch.</p> <p>Avoid exposure to extreme heat and cold.</p>
---	---

WARNING: This product is intended to be operated **ONLY** from the voltages listed on the back panel or the recommended, or included, power supply of the product. Operation from other voltages other than those indicated may cause irreversible damage to the product and void the products warranty. The use of AC Plug Adapters is cautioned because it can allow the product to be plugged into voltages in which the product was not designed to operate. If the product is equipped with a detachable power cord, use only the type provided with your product or by your local distributor and/or retailer. If you are unsure of the correct operational voltage, please contact your local distributor and/or retailer.

RACK MOUNTING

- A) Elevated Operating Ambient - If installed in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may be greater than room ambient. Therefore, consideration should be given to installing the equipment in an environment compatible with the maximum ambient temperature (T_{ma}) specified by the manufacturer.
- B) Reduced Air Flow - Installation of the equipment in a rack should be such that the amount of air flow required for safe operation of the equipment is not compromised.
- C) Mechanical Loading - Mounting of the equipment in the rack should be such that a hazardous condition is not achieved due to uneven mechanical loading.
- D) Circuit Overloading - Consideration should be given to the connection of the equipment to the supply circuit and the effect that overloading of the circuits might have on overcurrent protection and supply wiring. Appropriate consideration of equipment nameplate ratings should be used when addressing this concern.
- E) Reliable Earthing - Reliable earthing of rack-mounted equipment should be maintained. Particular attention should be given to supply connections other than direct connections to the branch circuit (e.g. use of power strips)."

FCC AND CANADA EMC COMPLIANCE INFORMATION:

This device complies with part 15 of the FCC Rules. Operation is subject to the following two conditions:

- (1) This device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

NOTE: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy, and if it is not installed and used in accordance with the instruction manual, it may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

CAN ICES-3 (A)/NMB-3(A)

EU COMPLIANCE INFORMATION:

Eligible to bear the CE mark; Conforms to European Union Low Voltage Directive 2014/35/EU; European Union EMC Directive 2004/108/EC; European Union Restriction of Hazardous Substances Recast (RoHS2) Directive 2011/65/EU; European Union WEEE (recast) Directive 2012/19/EU; European Union Eco-Design Directive 2009/125/EC; European Union Registration, Evaluation, Authorization and Restriction of Chemicals (REACH) Directive 2006/121/EC

You may obtain a free copy of the Declaration of Conformity by visiting <http://www.amx.com/techcenter/certifications.asp>.

WEEE NOTICE:

	<p>This appliance is labeled in accordance with European Directive 2012/19/EU concerning waste of electrical and electronic equipment (WEEE). This label indicates that this product should not be disposed of with household waste. It should be deposited at an appropriate facility to enable recovery and recycling.</p>
---	--

CHINA COMPLIANCE INFORMATION:



	<p>This device is designed and evaluated under the condition of non-tropical climate; it can only be used in locations in non-tropical climate areas. Using the device in tropical climate areas could result in a potential safety hazard.</p>
	<p>This device is designed and evaluated under the condition of altitude below 2000 meters above sea level; it can only be used in locations below 2000 meters above sea level. Using the device above 2000 meters could result in a potential safety hazard.</p>

Table of Contents

Overview	16
NXA-ENET8-POE+	16
10/100/1000BASE-T RJ-45 Ports	16
Port Status LEDs	16
Console Port	16
Gigabit SFP Slots.....	16
PoE Button	16
System LEDs.....	16
Factory Default Button.....	16
Cooling Fans and Vents	17
AC Power Socket.....	17
Grounding Terminal.....	17
Hardware Specifications	17
Web Interface	18
Key Features	18
Configuration Backup and Restore	18
Authentication	18
General Security Measures	18
Access Control Lists	18
DHCP/DHCPv6	18
Port Configuration	18
Port Mirroring	18
Port Trunking.....	18
Congestion Control.....	18
Address Table.....	18
IP Version 4 and 6	18
IEEE 802.1D Bridge	18
Store-and-Forward Switching	18
Spanning Tree Algorithm	18
Virtual LANs	18
Traffic Prioritization.....	18
Quality of Service.....	18
Link Layer Discovery Protocol	18
Multicast Filtering.....	18
Description of Software Features	19
Configuration Backup and Restore.....	19
Authentication	19
Access Control Lists.....	19
Port Configuration	19
Rate Limiting	19
Port Mirroring	19
Port Trunking.....	19
Storm Control	19
Static MAC Addresses.....	19
IP Address Filtering	19
IEEE 802.1D Bridge.....	19
Store-and-Forward Switching	20
Spanning Tree Algorithm.....	20
Virtual LANs	20
Traffic Prioritization.....	20

Quality of Service	20
Address Resolution Protocol	20
Multicast Filtering	20
Link Layer Discovery Protocol	21
System Defaults	21
Console Port Connection	21
Authentication and Security Measures	21
Web Management	21
SNMP	21
Port Configuration	21
Port Trunking	21
Congestion Control	22
Address Table	22
Spanning Tree Algorithm	22
Virtual LANs	22
Traffic Prioritization	22
IP Settings	22
System Log	22
SMTP Email Alerts	22
SNTP	22
Installation	23
Package Contents	23
Optional Item	23
Switch Installation Tasks	23
Unpack Package and Check Contents	23
Install the Chassis	23
Connect AC Power to Power On	24
Verify Switch Operation	24
Make Initial Configuration Changes	24
Through an RJ-45 Port	24
Through the Console Port	25
Install Transceivers and Connect Cables	25
Switch Chassis	25
General Installation Guidelines	25
Switch Cooling Requirements	26
Rack Cooling	26
How to Install the Switch in a Rack	26
Rack-Mounting the Switch	27
Installing the Switch on a Shelf or Desktop	27
Power and Grounding	28
Switch Power Supply	28
Grounding the Chassis	28
Connecting to AC Power	28
Port Connections	29
Cable Labeling and Connection Records	29
Understanding the Port Status LEDs	29
Installing an SFP Transceiver	29
Connecting to Twisted-Pair Copper Ports	30
Copper Cabling Guidelines	30

10/100BASE-TX Pin Assignments	31
1000BASE-T Assignments.....	31
1000BASE-T Cable Requirements.....	31
Power-over-Ethernet	32
Connecting to SFP Fiber Optic Ports.....	32
Switch Management	34
Understanding the System Status LEDs	34
Connecting to the Console Port.....	34
Resetting the Switch.....	36
Resetting to the Saved Configuration File.....	36
Resetting to the Factory Default Settings.....	36
Using the Web Interface	37
Overview	37
Connecting to the Web Interface	37
Navigating the Web Browser Interface	37
Dashboard.....	37
Home Page.....	38
Configuration Options	38
Panel Display	38
Basic Management Tasks	39
Overview	39
System Information.....	39
Hardware/Software Versions.....	40
System Capabilities	40
Configuring Support for Jumbo Frames	41
Usage Guidelines.....	41
Displaying Bridge Extension Capabilities	41
System Files.....	42
Copying Files via FTP/TFTP or HTTP	42
Command Usage.....	42
Saving the Running Configuration to a Local File.....	43
Setting the Start-up File	44
Showing System Files	44
Automatic Operation Code Upgrade	45
Usage Guidelines.....	45
Setting the System Clock.....	47
Setting the Time Manually	47
Setting the SNTP Polling Interval	48
Configuring NTP.....	49
Configuring Time Servers	49
Specifying SNTP Time Servers	49
Specifying NTP Time Servers.....	50
Specifying NTP Authentication Keys	51
Setting the Time Zone.....	52

Configuring Summer Time.....	52
Configuring the Console Port	54
Configuring Telnet Settings	55
Displaying CPU Utilization	56
Configuring CPU Guard	56
Displaying Memory Utilization.....	57
Resetting the System	58
Command Usage.....	58
Interface Configuration	61
Overview	61
Port Configuration.....	61
Configuring by Port List	61
Command Usage.....	61
Configuring by Port Range	62
Displaying Connection Status.....	63
Showing Port or Trunk Statistics.....	63
Displaying Statistical History.....	66
Command Usage.....	66
Displaying Transceiver Data	69
Configuring Transceiver Thresholds	69
Trunk Configuration	71
Command Usage.....	71
Configuring a Static Trunk.....	71
Command Usage.....	71
Configuring a Dynamic Trunk	73
Command Usage.....	73
Displaying LACP Port Counters.....	76
Displaying LACP Settings and Status for the Local Side.....	77
Displaying LACP Settings and Status for the Remote Side.....	78
Configuring Load Balancing	79
Command Usage.....	79
Saving Power	80
Command Usage.....	80
Configuring Local Port Mirroring.....	81
Command Usage.....	81
Configuring Remote Port Mirroring.....	82
Command Usage.....	82
Configuration Guidelines.....	82
RSPAN Limitations	82
Traffic Segmentation	84
Enabling Traffic Segmentation	84
Configuring Uplink and Downlink Ports.....	85
Command Usage.....	85

VLAN Configuration	87
IEEE 802.1Q VLANs.....	87
Assigning Ports to VLANs	87
Configuring VLAN Groups	88
Adding Static Members to VLANs	89
Protocol VLANs	92
Command Usage.....	92
Configuring Protocol VLAN Groups	92
Mapping Protocol Groups to Interfaces	93
Command Usage.....	93
Configuring MAC-based VLANs	94
Command Usage.....	94
Address Table Settings	96
Configuring MAC Address Learning	96
Command Usage.....	96
Setting Static Addresses	97
Command Usage.....	97
Changing the Aging Time	98
Displaying the Dynamic Address Table	98
Clearing the Dynamic Address Table	99
Issuing MAC Address Traps	100
Spanning Tree Algorithm	101
Overview	101
Configuring Loopback Detection.....	103
Configuring Global Settings for STA	104
Command Usage.....	104
Displaying Global Settings for STA.....	107
Configuring Interface Settings for STA.....	108
Displaying Interface Settings for STA	111
Configuring Multiple Spanning Trees.....	113
Command Usage.....	113
Configuring Interface Settings for MSTP	115
Congestion Control	117
Rate Limiting.....	117
Storm Control	118
Command Usage.....	118
Class of Service	119
Layer 2 Queue Settings	119
Setting the Default Priority for Interfaces	119
Command Usage.....	119

Selecting the Queue Mode	120
Command Usage.....	120
Layer 3/4 Priority Settings	121
Setting Priority Processing to DSCP or CoS	121
Command Usage.....	121
Mapping Ingress DSCP Values to Internal DSCP Values	122
Command Usage.....	122
Mapping CoS Priorities to Internal DSCP Values.....	123
Command Usage.....	123
Quality of Service	125
Overview	125
Command Usage.....	125
Configuring a Class Map.....	125
Command Usage.....	125
Creating QoS Policies	128
Attaching a Policy Map to a Port	130
Command Usage.....	130
VoIP Traffic Configuration	131
Overview	131
Configuring VoIP Traffic.....	131
Command Usage.....	131
Configuring Telephony OUI.....	132
Configuring VoIP Traffic Ports	133
Command Usage.....	133
Security Measures	134
AAA (Authentication, Authorization, and Accounting)	134
Configuring Local/ Remote Logon Authentication.....	135
Command Usage.....	135
Configuring Remote Login Authentication Servers.....	135
Command Usage.....	136
Configuring AAA Accounting	138
Command Usage.....	138
Configuring AAA Authorization	142
Command Usage.....	142
Configuring User Accounts	144
Command Usage.....	144
Network Access (MAC Address Authentication).....	145
Command Usage.....	145
Configuring Global Settings for Network Access	146
Configuring Network Access for Ports	147
Configuring a MAC Address Filter	148
Command Usage.....	148
Displaying Secure MAC Address Information	149

Configuring HTTPS	150
Configuring Global Settings for HTTPS	150
Command Usage.....	150
Replacing the Default Secure-site Certificate	151
Configuring the Secure Shell	151
Command Usage.....	152
Configuring the SSH Server	153
Generating the Host Key Pair	153
Importing User Public Keys	154
Access Control Lists	156
Command Usage.....	156
Showing TCAM Utilization.....	156
Command Usage.....	156
Setting the ACL Name and Type	157
Configuring a Standard IPv4 ACL.....	158
Configuring an Extended IPv4 ACL	159
Configuring a Standard IPv6 ACL.....	160
Configuring an Extended IPv6 ACL.....	161
Configuring a MAC ACL.....	163
Configuring an ARP ACL.....	164
Binding a Port to an Access Control List.....	165
Showing ACL Hardware Counters	166
ARP Inspection	167
Command Usage.....	167
Configuring Global Settings for ARP Inspection	167
Command Usage.....	167
Configuring VLAN Settings for ARP Inspection	168
Command Usage.....	168
Configuring Interface Settings for ARP Inspection.....	169
Displaying ARP Inspection Statistics.....	170
Displaying the ARP Inspection Log.....	170
Filtering IP Addresses for Management Access	171
Command Usage.....	171
Configuring Port Security	172
Command Usage.....	172
Configuring 802.1x Port Authentication	173
Configuring 802.1x Global Settings	174
Configuring Port Authenticator Settings for 802.1x	174
Command Usage.....	174
Displaying 802.1x Statistics	176
DHCP Snooping	177
Command Usage.....	177
DHCP Snooping Global Configuration	178

DHCP Snooping VLAN Configuration	179
Command Usage.....	179
Configuring Ports for DHCP Snooping	180
Command Usage.....	180
Displaying DHCP Snooping Binding Information	181
DoS Protection.....	182
IPv4 Source Guard	183
Configuring Ports for IPv4 Source Guard.....	183
Command Usage.....	183
Configuring Static Bindings for IPv4 Source Guard.....	184
Command Usage.....	184
Displaying Information for Dynamic IPv4 Source Guard Bindings	186
Basic Administration Protocols	187
Configuring Event Logging	187
System Log Configuration	187
Remote Log Configuration.....	188
Sending Simple Mail Transfer Protocol Alerts.....	189
Link Layer Discovery Protocol	190
Setting LLDP Timing Attributes.....	190
Configuring LLDP Interface Attributes	191
Configuring LLDP Interface Civic-Address	194
Command Usage.....	194
Displaying LLDP Local Device Information	195
Displaying LLDP Remote Device Information	197
Displaying Device Statistics.....	202
Power over Ethernet.....	203
Setting the Switch's Overall PoE Power Budget	204
Setting the Port PoE Power Budget	205
Command Usage.....	205
Simple Network Management Protocol.....	206
Command Usage.....	207
Configuring Global Settings for SNMP.....	207
Setting the Local Engine ID	208
Command Usage.....	208
Specifying a Remote Engine ID.....	208
Command Usage.....	208
Setting SNMPv3 Views	209
Configuring SNMPv3 Groups	211
Setting Community Access Strings	214
Configuring Local SNMPv3 Users	215
Configuring Remote SNMPv3 Users	217
Command Usage.....	217
Specifying Trap Managers.....	218
Command Usage.....	218

Creating SNMP Notification Logs.....	221
Command Usage.....	221
Showing SNMP Statistics	222
Remote Monitoring	223
Configuring RMON Alarms	223
Command Usage.....	223
Configuring RMON Events.....	225
Command Usage.....	225
Configuring RMON History Samples	226
Command Usage.....	226
Configuring RMON Statistical Samples.....	228
Command Usage.....	228
Setting a Time Range	229
Command Usage.....	229
LBD Configuration	231
Usage Guidelines	231
Configuring Global Settings for LBD	231
Configuring Interface Settings for LBD	232
Multicast Filtering	233
Overview	233
Layer 2 IGMP (Snooping and Query for IPv4).....	233
Configuring IGMP Snooping and Query Parameters.....	234
Command Usage.....	234
Specifying Static Interfaces for a Multicast Router.....	236
Command Usage.....	236
Assigning Interfaces to Multicast Services	238
Command Usage.....	238
Setting IGMP Snooping Status per Interface.....	239
Command Usage.....	239
Filtering IGMP Query Packets and Multicast Data	242
Displaying Multicast Groups Discovered by IGMP Snooping	243
Command Usage.....	243
Displaying IGMP Snooping Statistics	243
Filtering and Throttling IGMP Groups	246
Enabling IGMP Filtering and Throttling	246
Configuring IGMP Filter Profiles.....	246
Command Usage.....	246
Configuring IGMP Filtering and Throttling for Interfaces.....	248
Command Usage.....	248
MLD Snooping (Snooping and Query for IPv4)	249
Configuring MLD Snooping and Query Parameters.....	249
Setting Immediate Leave Status for MLD Snooping per Interface.....	250
Specifying Static Interfaces for an IPv6 Multicast Router	250
Command Usage.....	250
Assigning Interfaces to IPv6 Multicast Services.....	252

Command Usage.....	252
Showing MLD Snooping Groups and Source List.....	253
IP Tools	254
Using the Ping Function	254
Command Usage.....	254
Using the Trace Route Function	255
Command Usage.....	255
Address Resolution Protocol	256
Displaying Dynamic or Local ARP Entries	256
IP Services	257
Domain Name Service.....	257
Configuring General DNS Service Parameters	257
Command Usage.....	257
Configuring a List of Domain Names	257
Command Usage.....	257
Configuring a List of Name Servers	258
Command Usage.....	258
Configuring Static DNS Host to Address Entries	259
Command Usage.....	259
Displaying the DNS Cache.....	260
Command Usage.....	260
Dynamic Host Configuration Protocol	260
Specifying a DHCP Client Identifier	260
Command Usage.....	260
Configuring DHCP Relay Service.....	261
Command Usage.....	262
Enabling DHCP Dynamic Provision	263
Command Usage.....	263
IP Configuration	265
Setting the Switch's IP Address (IP Version 4)	265
Configuring the IPv4 Default Gateway	265
Configuring IPv4 Interface Settings	265
Setting the Switch's IP Address (IP Version 6)	267
Command Usage.....	267
Configuring the IPv6 Default Gateway	267
Configuring IPv6 Interface Settings	268
Command Usage.....	268
Configuring an IPv6 Address.....	270
Command Usage.....	270
Showing IPv6 Addresses	272
Showing the IPv6 Neighbor Cache	273
Showing IPv6 Statistics	274
Command Usage.....	274
Showing the MTU for Responding Destinations	277

Appendix A: Software Specifications	278
Software Features	278
Management Authentication	278
General Security Measures.....	278
Port Configuration	278
Flow Control.....	278
Storm Control	278
Port Mirroring	278
Rate Limits	278
Port Trunking.....	278
Spanning Tree Algorithm.....	278
VLAN Support.....	278
Class of Service.....	278
Quality of Service	278
Multicast Filtering	278
IP Routing	278
Additional Features	278
Management Features.....	279
In-Band Management.....	279
Out-of-Band Management	279
Software Loading.....	279
SNMP.....	279
RMON	279
Standards	279
Management Information Bases.....	279
Appendix B: Troubleshooting	281
Diagnosing LED Indicators	281
System Self-Diagnostic Test Failure.....	281
Power and Cooling Problems	281
Installation	281
In-Band Access	281
Problems Accessing the Management Interface.....	282
Using System Logs	282

Overview

NXA-ENET8-POE+

The NXA-ENET8-POE+ (FG2178-64) is a Gigabit Ethernet switch with 8 10/100/1000BASE-T ports, and two Small Form Factor Pluggable (SFP) transceiver slots for fiber connectivity. The switch includes an SNMP-based management agent, which provides both in-band and out-of-band access for managing the switch. Further, the switches support both web and CLI-based configuration.

All of the 10/100/1000 Mbps ports on the NXA-ENET-POE+ support both the IEEE 802.3af-2003 and IEEE 802.3at-2009 PoE standards. The switch is an excellent choice for supplying power to connected PoE devices such as web cameras, IP telephones, or access points.

The switch consists of several key hardware components. This manual describes each specific component, or related components, together with their installation requirements and procedures in each chapter. To understand each component in detail, refer to the relevant section.



FIG. 1 NXA-ENET-POE+ (front and rear panels)

10/100/1000BASE-T RJ-45 Ports

The switch contains 8 10/100/1000BASE-T RJ-45 ports that support 10/100/ 1000BASE-T copper links to other devices. For more information, see the *Connecting to Twisted-Pair Copper Ports* section on page 30.

Port Status LEDs

For information on port status LED indicators, see the *Understanding the Port Status LEDs* section on page 29.

Console Port

The RJ-45 connector on the front panel right side that is labeled "Console" provides an out-of-band serial connection to a terminal or a PC running terminal emulation software. The port can be used for performing switch monitoring and configuration. For more information, see the *Connecting to the Console Port* section on page 34.

Gigabit SFP Slots

The switches contain up to four SFP transceiver slots that operate up to 1 Gbps full duplex. For more information, see the *Connecting to SFP Fiber Optic Ports* section on page 32.

PoE Button

Pressing the PoE button on the front panel changes the port LEDs to display PoE status. For more information, see the *Understanding the System Status LEDs* section on page 34.

System LEDs

For information on system status LED indicators, see the *Understanding the System Status LEDs* section on page 34.

Factory Default Button

Pressing the reset button on the front panel causes the switch to restart or restore factory default settings. For more information, see the *Resetting the Switch* section on page 36.

Cooling Fans and Vents

The switch must be installed in a properly cooled and ventilated environment. For more information, see the *Rack Cooling* section on page 26.

AC Power Socket

The switch requires a 100-240 VAC, 50-60 Hz AC power source. For more information on the switch power input, how to connect it, and how to power-on the switch, see the *Connecting to AC Power* section on page 28.

Grounding Terminal

The switch includes a grounding terminal that must be connected to a ground source that provides local earth potential. For more information, see the *Grounding the Chassis* section on page 28.

Hardware Specifications

NXA-ENET8-POE+ Hardware Specifications	
Ports	8 1000BASE-T RJ-45 ports with Auto-negotiation, 2 Gigabit SFP transceiver slots
Network Interface	<ul style="list-style-type: none"> Ports 1-8: RJ-45 connector, auto MDI/X Ports 9-10: Gigabit SFP transceivers Packet Buffer Size: 1 MB
Buffer Architecture	4 Mbytes
Aggregate Bandwidth	104 Gbps
Switching Database	8K MAC address entries
AC Input Power	AC 100-240 V, 50-60 Hz, 2.1 A
Power Consumption	160W
Weight	2.4kg (5.34lb)
Size	33.0 x 20.4 x 4.26 cm (12.99 x 8.03 x 1.67 in.)
Temperature	<ul style="list-style-type: none"> Operating: 0° C to 40° C (32° F to 104° F) Storage: -40° C to 70° C (-40° F to 158° F)
Humidity	Operating: 10% to 90% (non-condensing)
Out-of-Band Management	Front panel RJ-45 console port
In-Band Management	SSH, Telnet, SNMP, or HTTP
Software Loading	HTTP, FTP/TFTP in-band
Forwarding Mode	Store-and-forward
Throughput	Wire speed
Flow Control	<ul style="list-style-type: none"> Full Duplex: IEEE 802.3x Half Duplex: Back pressure

Web Interface

The NXA-ENET8-POE+ provides a broad range of features for Layer 2 switching and Layer 3 routing. It includes a management agent that allows you to configure the features listed in this manual. The default configuration can be used for most of the features provided by this switch. However, there are many options that you should configure to maximize the switch's performance for your particular network environment.

Key Features

The NXA-ENET8-POE+ provides a wide range of advanced performance enhancing features, as described below:

NXA-ENET8-POE+ Key Features	
Configuration Backup and Restore	Using management station or TFTP server
Authentication	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • AAA • ARP Inspection • DHCP Snooping (with Option 82 relay information) • DoS Protection • IP Source Guard • Port Authentication - IEEE 802.1X • Port Security - MAC address filtering
Access Control Lists	Supports up to 256 ACLs, 128 rules per ACL, and 512 rules per system
DHCP/DHCPv6	Client, Relay, Relay Option 82
Port Configuration	Speed, duplex mode, and flow control
Port Mirroring	3 sessions, one or more source ports to an analysis port
Port Trunking	Supports up to 8 trunks - static or dynamic trunking (LACP)
Congestion Control	<ul style="list-style-type: none"> • Rate Limiting • Throttling for broadcast, multicast, unknown unicast storms
Address Table	<ul style="list-style-type: none"> • 8K MAC addresses in the forwarding table (shared with L2 unicast, L2 multicast, IPv4 multicast, IPv6 multicast) • 1K static MAC addresses • 512 L2 IPv4 multicast groups (shared with MAC address table)
IP Version 4 and 6	Supports IPv4 and IPv6 addressing and management
IEEE 802.1D Bridge	Supports dynamic data switching and addresses learning
Store-and-Forward Switching	Supported to ensure wire-speed switching while eliminating bad frames
Spanning Tree Algorithm	Supports standard STP, Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Trees (MSTP)
Virtual LANs	Up to 4094 using IEEE 802.1Q, port-based, protocol-based, voice VLANs, and QinQ tunnel
Traffic Prioritization	Default port priority, traffic class map, queue scheduling, IP Precedence, or Differentiated Services Code Point (DSCP)
Quality of Service	Supports Differentiated Services (DiffServ)
Link Layer Discovery Protocol	Used to discover basic information about neighboring devices
Multicast Filtering	Supports IGMP snooping and query for Layer 2

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 TFTP server (using the web or console interface), and later download this file to restore the switch configuration settings.

Authentication

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

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

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 traffic passing through the port is restricted. If traffic rises above a pre-defined threshold, it will be throttled until the level falls back beneath the threshold.

Static MAC Addresses

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

IP Address Filtering

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

IEEE 802.1D Bridge

The switch supports IEEE 802.1D transparent bridging. The address table facilitates data switching by learning addresses, and then filtering or forwarding traffic based on this information. The address table supports up to 16K addresses.

Store-and-Forward Switching

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

To avoid dropping frames on congested ports, the switch provides 12 Mbits for frame buffering. This buffer can queue packets awaiting transmission on congested networks.

Spanning Tree Algorithm

The switch supports these spanning tree protocols:

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

Virtual LANs

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

- Eliminate broadcast storms which severely degrade performance in a flat network.
- Simplify network management for node changes/moves by remotely configuring VLAN membership for any port, rather than having to manually change the network connection.
- Provide data security by restricting all traffic to the originating VLAN, except where a connection is explicitly defined via the switch's routing service.
- Use private VLANs to restrict traffic to pass only between data ports and the uplink ports, thereby isolating adjacent ports within the same VLAN, and allowing you to limit the total number of VLANs that need to be configured.
- Use protocol VLANs to restrict traffic to specified interfaces based on protocol type.

Traffic Prioritization

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

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

Quality of Service

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

Address Resolution Protocol

The switch uses ARP to convert between IP addresses and MAC (hardware) addresses. This switch supports conventional ARP, which locates the MAC address corresponding to a given IP address. 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.

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.

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.

NXA-ENET8-POE+ System Defaults		
Function	Parameter	Default
Console Port Connection	Baud Rate	115200 bps
	Data bits	8
	Stop bits	1
	Parity	none
	Local Console Timeout	600 seconds
Authentication and Security Measures	Privileged Exec Level	Username: admin Password: admin
	Normal Exec Level	Username: guest Password: guest
	Enable Privileged Exec from Normal Exec Level	Password: super
	RADIUS Authentication	Disabled
	TACACS+ Authentication	Disabled
	802.1X Port Authentication	Disabled
	MAC Authentication	Disabled
	HTTPS	Enabled
	SSH	Disabled
	Port Security	Disabled
	IP Filtering	Disabled
	DHCP Snooping	Disabled
	IP Source Guard	Disabled (all ports)
Web Management	HTTP Server	Enabled
	HTTP Port Number	80
	HTTP Secure Server	Enabled
	HTTP Secure Server Port	443
SNMP	SNMP Agent	Enabled
	Community Strings	"public" (read only) "private" (read/write)
	Traps	Authentication traps: enabled Link-up-down events: enabled
	SNMP V3	• View: defaultview • Group: public (read-only); private (read/write)
Port Configuration	Admin Status	Enabled
	Auto-negotiation	Enabled
	Flow Control	Disabled
Port Trunking	Static Trunks	None
	LACP (all ports)	Disabled

NXA-ENET8-POE+ System Defaults (Cont.)		
Function	Parameter	Default
Congestion Control	Rate Limiting	Disabled
	Storm Control	Broadcast: Enabled (64kbits/sec) 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)	Hybrid
	GVRP (global)	Disabled
	GVRP (port interface)	Disabled
	QinQ Tunneling	Disabled
Traffic Prioritization	Ingress Port Priority	0
	Queue Mode	WRR
	Queue Weight	Queue: 0 1 2 3 4 5 6 7 Weight: 1 2 4 6 8 10 12 14
	Class of Service	Enabled
	IP DSCP Priority	Disabled
IP Settings	Management. VLAN	VLAN 1
	IP Address	DHCP assigned
	Subnet Mask	255.255.255.0
	Default Gateway	Not configured
	DHCP	Client: Enabled
	BOOTP	Disabled
	ARP	Enabled Cache Timeout: 20 minutes
Multicast Filtering	IGMP Snooping (Layer 2)	Snooping: Enabled 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
SMTP Email Alerts	Event Handler	Enabled (but no server defined)
SNTP	Clock Synchronization	Disabled

Installation

This chapter features the following sections:

- Package Contents
- Switch Installation Tasks

Package Contents

After unpacking the switch, check the contents to be sure you have received all the components.

- NXA-ENET8-POE+ Gigabit PoE Ethernet Switch
- AC Power Cord
- Rack Mounting Kit containing two standard brackets and eight screws for attaching the brackets to the switch.
- 4 adhesive foot pads for surface mounting
- Quick Start Guide
- Safety Instructions

Optional Item

- Console cable-RJ-45 to DB-9

NOTE: Other documentation including the Quick Start Guide and CLI Reference Guide can be accessed on the NXA-ENET8-POE+ product page at www.amx.com.

Switch Installation Tasks

Follow these tasks to install the switch in your network. For full details on each task, go to the relevant chapter or section by clicking on the link.

CAUTION: Before installing your switch, first review all the safety statements and guidelines in the Regulatory and Safety Information document.

Unpack Package and Check Contents

Unpack your switch and check the package contents to be sure you have received all the items. See *Package Contents* section on page 23 for more information.

Install the Chassis

The switch is designed to be installed in a standard 19-inch equipment rack. Plan your rack installation and install the switch chassis in the rack. Be sure to take into account switch cooling requirements.

See the *Switch Chassis* section on page 25 for more information.



FIG. 2 Installing the Switch in a Rack

1. Attach the brackets to the switch.
2. Use the screws supplied with the rack to secure the switch in the rack.

Connect AC Power to Power On

Prior to connecting to AC power, assure to connect the chassis ground connection to a known earth ground. Connect the power cord to the AC socket on the switch and to a grounded, 3-pin, AC power source.

See the *Power and Grounding* section on page 28 for more information.



FIG. 3 Connecting AC Power

1. Connect a grounding wire to the grounding terminal.
2. Connect an external AC power source to the AC power socket of the switch using the supplied AC power cord.

Verify Switch Operation

Verify basic switch operation by checking the system LEDs.

When operating normally, the Power and Diag LED should both be on green. If either of these LEDs are on amber, see the *Diagnosing LED Indicators* section on page 281 for more information.

Go to the *Understanding the System Status LEDs* section on page 34 for more information.



FIG. 4 System Status LEDs

Make Initial Configuration Changes

At this point, you may need to make a few basic switch configuration changes before connecting to the network. You can either connect to the switch console port or any RJ-45 port to perform this task.

Through an RJ-45 Port

The switch offers a user-friendly web-based management interface for the configuration of all the unit's features.

You can make initial configuration changes by connecting a PC directly to one of the switch's RJ-45 ports. The switch has a default management IP address of 192.168.2.10 and a subnet mask of 255.255.255.0. You must set your PC IP address to be on the same subnet as the switch (that is, the PC and switch addresses must both start 192.168.2.x).

Log in to the web interface using the default settings:

- Login Name - admin
- Password - admin

Through the Console Port

The serial port's configuration requirements are as follows: 115200 bps, 8 characters, no parity, one stop bit, 8 data bits, and no flow control.

You can log in to the command-line interface (CLI) using default settings: User "admin" with password "admin".

See the *Connecting to the Console Port* section on page 34 for more information.



FIG. 5 Console Port

1. Connect console cable to switch's Console port.
2. Connect console cable to PC's DB-9 COM port.

For information on initial switch configuration, refer to the CLI Reference Guide.

Install Transceivers and Connect Cables

Install SFP transceivers and connect network cables to port interfaces:

- For RJ-45 ports, use 100-ohm category 3 or better Ethernet cable for 10BASE-T connections, use 100-ohm category 5 or better Ethernet cable for 100BASE-TX and 1000BASE-T connections.
- Install SFP transceivers and then connect fiber optic cabling to the transceiver ports.

As connections are made, check the port status LEDs to be sure the links are valid.

See the *Port Connections* section on page 29 for more information.




FIG. 6 RJ-45 Ports

1. 10/100/1000BASE-T RJ-45 Port.
2. Twisted-pair Cable with RJ-45 Plug.

Switch Chassis

The switch is designed to be installed in a standard 19-inch equipment rack.

General Installation Guidelines

Be sure to follow the guidelines below when choosing a location. The installation location should conform to the following criteria:

- Be able to maintain its temperature within 0 to 50°C/32 to 122°F (0 to 40°C/ 32 to 104°F) and its humidity within 10% to 90%, non- condensing.
- Provide adequate space (approximately five centimeters or two inches) on all sides for proper air flow.
- Be accessible for installing, cabling and maintaining the device.
- Allow the status LEDs to be clearly visible.
- Make sure twisted-pair cable is always routed away from power lines, fluorescent lighting fixtures and other sources of electrical interference, such as radios and transmitters.
- Make sure that the unit is connected to a separate grounded power outlet and is powered from an independent circuit breaker. As with any equipment, using a filter or surge suppressor is recommended. Verify that the external AC power requirements for the switch can be met as listed under the *Switch Power Supply* section on page 28.

Switch Cooling Requirements

Wherever the switch is located, be sure to pay close attention to switch cooling requirements. The location should be well ventilated and provide unrestricted airflow at the front, back, and sides of the switch. If the airflow is insufficient, it may cause the switch to overheat and possibly fail.

The NXA-ENET8-POE+ uses a fanless cooling design. The following figure shows the convective airflow from the switch.



FIG. 7 Switch Cooling

Rack Cooling

When mounting the switch in an enclosed rack or cabinet, be sure to check the following guidelines to prevent overheating:

- Make sure that enough cool air can flow into the enclosure for the equipment it contains.
- Check that the rack or cabinet allows the hot air to exit the enclosure (normally from the top) without circulating back into equipment.
- If the enclosure has sides or doors with ventilation holes, make sure they are not blocked by cables or other obstructions.
- Route cables within the rack or cabinet to maximize the airflow.
- When possible, do not completely fill the rack or cabinet with equipment, allow some unused space within the enclosure for better airflow.

How to Install the Switch in a Rack

When rack mounting the switch, pay particular attention to the following factors:

- **Rack Types:** You can use any standard EIA 19-inch equipment rack with either two or four posts. The bracket hole pattern should be spaced 1U (1.75 in. or 4.45 cm) apart.
- **Rack Stability:** Whenever possible, secure the rack to the building ceiling or floor, particularly if you are located in a region where earthquakes are common.
- **Rack Planning:** When installing equipment in a rack, first plan how units can be best arranged. Try to always mount the heaviest equipment at the bottom of the rack.
- **Temperature:** Since the temperature within a rack assembly may be higher than the ambient room temperature, check that the rack-environment temperature is within the specified operating temperature range. See the *Switch Cooling Requirements* section on page 26 for more information.
- **Mechanical Loading:** Do not place any equipment on top of a rack-mounted unit.
- **Circuit Overloading:** Be sure that the supply circuit to the rack assembly is not overloaded.
- **Grounding:** Rack-mounted equipment should be properly grounded.

Rack-Mounting the Switch

Before you start to rack-mount the switch, be sure to have the following items available:

- Four mounting screws for each device you plan to install in a rack—these are not included. Be sure to use the rack mounting screws that are supplied with the rack.
- A screwdriver (Phillips or flathead, depending on the type of screws used).

Perform the following steps to rack mount the switch:

CAUTION: *Installing the switch in a rack requires two people: One should position the switch in the rack, while the other secures it using the mounting screws.*

1. Attach the brackets to the device using the screws provided in the Rack Mounting Kit.



FIG. 8 Attaching the brackets

2. Following your rack plan, mark the holes in the rack where the switch will be installed.
3. One person should lift the switch into the rack so that it is aligned with the marked holes.
4. The second person should secure the switch in the rack, using four rack-mounting screws (not provided).




FIG. 9 Installing the switch in a rack

5. If installing a single switch only, go to the *Power and Grounding* section on page 28.
6. If installing multiple switches, repeat steps 1 to 4 to mount the switches according to your rack plan.

Installing the Switch on a Shelf or Desktop

The switch can be installed on any flat surface such as a desktop or shelf. To mount the switch on a flat surface, follow these steps:

1. Attach the four adhesive feet to the bottom of the switch.
2. Set the device on a flat surface near an AC power source, making sure there are at least two inches of space on all sides for proper airflow.
3. If installing a single switch only, go to the *Connecting to AC Power* section on page 28.
4. If installing multiple switches, attach four adhesive feet to each one. Place each device squarely on top of the one below, in any order.

Power and Grounding

The following sections provide details on how to connect AC power to the switch, grounding the chassis, and how to power-on the switch.

Switch Power Supply

The switch requires power from an external AC power supply that can meet the required specification described below.

Active power requirements:

- AC Input Power: 100-240V, 50-60Hz, 2.1A
- Total Power consumption: 160W

NOTE: *Maximum power consumption values are measured under a 100 percent loading test and should be used as estimates for planning purposes.*

A standard AC power socket is located on the rear panel of the switch. The power socket is for the AC power cord.



FIG. 10 AC Power Supply Socket

Grounding the Chassis

The rear panel of the switch chassis includes a single hole grounding terminal. It must be connected to ground to ensure proper operation and to meet electromagnetic interference (EMI) and safety requirements.



FIG. 11 Grounding Terminal

Before powering on the switch, ground the switch to earth as described below.

1. Ensure that the rack in which the switch is to be mounted is properly grounded and in compliance with ETSI ETS 300 253.
2. Ensure that there is a good electrical connection to the grounding point on the rack (no paint or isolating surface treatment).
3. Disconnect all power cables to the switch.
4. Attach a 6 AWG stranded copper wire to the grounding terminal on the switch.

The switch chassis is connected internally to 0 V. This circuit is connected to the single-hole grounding terminal on the rear panel of the switch (left of the AC power socket). The surface area around this terminal is not painted in order to provide for a good connection.

5. Attach the grounding wire to the ground point on the rack.

CAUTION: *The earth connection must not be removed unless all supply connections have been disconnected.*

Connecting to AC Power

Connect the switch to an AC power source to power on. Verify that the external AC power requirements for the switch can be met as listed below:

AC 100-240 V, 50-60 Hz, 2.1 A

To connect the switch to a power source:

1. Plug the power cord into a grounded, 3-pin, AC power source.
2. Insert the plug on the other end of the power cord directly into the AC input socket on the back of the switch (see FIG. 10).

NOTE: *If your country's AC power outlet standards do not match the power plug of the included AC power cord, you will need to change the AC power cord. You must use a cord set that has been approved for the socket type in your country.*

3. Check the LED indicators on the switch front panel as the unit is powered on to verify that power is being received. If not, recheck the power cord connections at the AC supply source and back panel power input connector.

Port Connections

This section provides details on making connections to switch network interfaces, including how to install optional transceivers, and details on network cable specifications.

Cable Labeling and Connection Records

When planning a network installation, it is essential to label the opposing ends of cables and to record where each cable is connected. Doing so will enable you to easily locate inter-connected devices, isolate faults and change your topology without need for unnecessary time consumption.

To best manage the physical implementations of your network, follow these guidelines:

- Clearly label the opposing ends of each cable.
- Using your building's floor plans, draw a map of the location of all network- connected equipment. For each piece of equipment, identify the devices to which it is connected.
- Note the length of each cable and the maximum cable length supported by the switch ports.
- For ease of understanding, use a location-based key when assigning prefixes to your cable labeling.
- Use sequential numbers for cables that originate from the same equipment.
- Differentiate between racks by naming accordingly.
- Label each separate piece of equipment.
- Display a copy of your equipment map, including meanings of all abbreviations at each equipment rack.

Understanding the Port Status LEDs

The switch includes LED indicators for each port to indicate link status and network activity.




FIG. 12 Port Status LEDs

The port LEDs are shown below and described in the following table.

Port Status LEDs		
1000BASE-T RJ-45 Ports 1-8		
LED	Condition	Status
Link/Activity	On/Blinking Amber	Port has a valid 10/ 100 Mbps link. Blinking indicates traffic on the port.
	On/Blinking Green	Port has a valid 1000 Mbps link. Blinking indicates traffic on the port.
	Off	The link is down.
PoE Mode*	On Amber	A PoE device is connected and delivering PoE power.
	Off	Not delivering PoE power.
Gigabit SFP Ports 9-10		
Link/Activity	On/Blinking Amber	Port has a valid 1000 Mbps link. Blinking indicates traffic on the port.
	On/Blinking Green	Port has a valid 1000 Mbps link. Blinking indicates traffic on the port.
	Off	The link is down.

* PoE Mode button is pressed

Installing an SFP Transceiver

The switch provides slots for optional SFP transceivers. The supported transceiver types are listed below:

- 1000BASE-SX
- 1000BASE-LX
- 1000BASE-ZX
- 1000BASE-BX10
- 1000BASE-BX20
- 100BASE-BX20

NOTE: SFP transceivers are hot-swappable. The switch does not need to be powered off before installing or removing a transceiver.

NOTE: SFP transceivers are not provided in the switch package.

Perform the following steps to install an SFP transceiver:

1. Consider network and cabling requirements to select an appropriate transceiver type that is also compatible with the switch transceiver support.
2. If the SFP slot is covered with a rubber protective cap, remove the cap and keep it for later replacement.
3. Insert the transceiver with the optical connector facing outward and the slot connector facing down. Note that SFP transceivers are keyed so they can only be installed in the correct orientation.
4. Slide the transceiver into the slot until it clicks into place. If you do not immediately connect a cable to the port, use a rubber protective cap to keep the transceiver optics clean.



FIG. 13 Inserting an SFP Transceiver into a Slot

NOTE: To remove a transceiver, first disconnect the network cable, then pull the tab to remove the transceiver from the slot.

Connecting to Twisted-Pair Copper Ports

The RJ-45 ports on the switch support automatic MDI/MDI-X pin-out configuration, which enables you to use standard straight-through twisted-pair cables to connect to any other network device (PCs, servers, switches, routers, or hubs).

The connection requires an unshielded twisted-pair (UTP) or shielded twisted-pair (STP) cable with RJ-45 connectors at both ends..

Maximum Twisted-Pair Copper Cable Lengths		
Cable Type	Maximum Cable Length	Connector
1000BASE-T Category 5, 5e, or 6 100-ohm UTP or STP	100m(328ft)	RJ-45
100BASE-TX Category 5 or better 100-ohm UTP or STP	100m(328ft)	RJ-45
10BASE-T Category 3 or better 100-ohm UTP	100m(328ft)	RJ-45

Copper Cabling Guidelines

To ensure proper operation when installing the switch into a network, make sure that the current cables are suitable for 10BASE-T, 100BASE-TX, or 1000BASE-T operation. Check the following criteria against the current installation of your network:

- Cable type: Unshielded twisted pair (UTP) or shielded twisted pair (STP) cables with RJ-45 connectors; Category 5, 5e or better cable for 1000BASE-T connections, Category 5 or better for 100BASE-TX connections, and Category 3 or better for 10BASE-T connections.
- Protection from radio frequency interference emissions
- Electrical surge suppression
- Separation of electrical wires (switch related or other) and electromagnetic fields from data based network wiring
- Safe connections with no damaged cables, connectors or shields

10/100BASE-TX Pin Assignments

All 100BASE-TX RJ-45 ports support automatic MDI/MDI-X operation, so you can use straight-through or crossover cables for all network connections to PCs, switches, or hubs. In straight-through cable, pins 1, 2, 3, and 6, at one end of the cable, are connected straight through to pins 1, 2, 3, and 6 at the other end of the cable.



FIG. 14 RJ-45 Connector

10/100BASE-TX MDI and MDI-X Port Pinouts		
Pin	MDI Signal Name*	MDI-X Signal Name
1	Transmit Data plus (TD+) -52V power (Negative Vport)	Receive Data plus (RD+) GND (Positive Vport)
2	Transmit Data minus (TD-) -52V power (Negative Vport)	Receive Data minus (RD-) GND (Positive Vport)
3	Receive Data plus (RD+) GND (Positive Vport)	Transmit Data plus (TD+) -52V power (Negative Vport)
4	-52V power (Negative Vport)	GND (Positive Vport)
5	-52V power (Negative Vport)	GND (Positive Vport)
6	Receive Data minus (RD-) GND (Positive Vport)	Transmit Data minus (TD-) -52V power (Negative Vport)
7	GND (Positive Vport)	-52V power (Negative Vport)
8	GND (Positive Vport)	-52V power (Negative Vport)

* The "+" and "-" signs represent the polarity of the wires that make up each wire pair.

1000BASE-T Assignments

All 1000BASE-T ports support automatic MDI/MDI-X operation, so you can use straight-through cables for all network connections to PCs, servers, or switches.

The table below shows the 1000BASE-T MDI and MDI-X port pinouts. These ports require that all four pairs of wires be connected. Note that for 1000BASE-T operation, all four pairs of wires are used for both transmit and receive.

1000BASE-T MDI and MDI-X Port Pinouts		
Pin	MDI Signal Name	MDI-X Signal Name
1	Bi-directional Pair A Plus (BI_DA+) -52V power (Negative Vport)	Bi-directional Pair B Plus (BI_DB+) GND (Positive Vport)
2	Bi-directional Pair A Minus (BI_DA-) -52V power (Negative Vport)	Bi-directional Pair B Minus (BI_DB-) GND (Positive Vport)
3	Bi-directional Pair B Plus (BI_DB+) GND (Positive Vport)	Bi-directional Pair A Plus (BI_DA+) -52V power (Negative Vport)
4	Bi-directional Pair C Plus (BI_DC+) -52V power (Negative Vport)	Bi-directional Pair D Plus (BI_DD+)- GND (Positive Vport)
5	Bi-directional Pair C Minus (BI_DC-) -52V power (Negative Vport)	Bi-directional Pair D Minus (BI_DD-) GND (Positive Vport)
6	Bi-directional Pair B Minus (BI_DB-) GND (Positive Vport)	Bi-directional Pair A Minus (BI_DA-) -52V power (Negative Vport)
7	Bi-directional Pair D Plus (BI_DD+) GND (Positive Vport)	Bi-directional Pair C Plus (BI_DC+) -52V power (Negative Vport)
8	Bi-directional Pair D Minus (BI_DD-) GND (Positive Vport)	Bi-directional Pair C Minus (BI_DC-) -52V power (Negative Vport)

1000BASE-T Cable Requirements

All Category 5 UTP cables that are used for 100BASE-TX connections should also work for 1000BASE-T, providing that all four wire pairs are connected. However, it is recommended that for all critical connections, or any new cable installations, Category 5e (enhanced Category 5) or Category 6 cable should be used. The Category 5e and 6 specifications include test parameters that are only recommendations for Category 5. Therefore, the first step in preparing existing Category 5 cabling for running 1000BASE-T is a simple test of the cable installation to be sure that it complies with the IEEE 802.3-2008 standards.

Power-over-Ethernet

The PoE switch supports both IEEE 802.3af and IEEE 802.3at-2009 PoE standards. These switches are excellent choices for supplying power to connected PoE devices such as web cameras, IP telephones, or access points.

NXA-ENET8-POE+ PoE Power Budget	
Total PoE Power Budget	125W
Ports supply up to 15.4W simultaneously	8
Ports supply up to 30W simultaneously	4

Any PoE-compliant device attached to a port can directly draw power from the switch over the Ethernet cable without requiring its own separate power source. This capability gives network administrators centralized power control for devices such as IP phones and wireless access points, which translates into greater network availability.

For each attached PoE-compliant device, the switch automatically senses the load and dynamically supplies the required power. The switch delivers power to a device using the wire pairs in UTP or STP cable.

Follow these steps to connect cables to 1000BASE-T RJ-45 twisted-pair copper ports:

1. Attach one end of a twisted-pair cable segment to the device's RJ-45 connector.
2. Attach the other end to an available port on the switch. Make sure each twisted pair cable does not exceed 100 meters (328 ft) in length.
3. As each connection is made, the Link LED (on the switch) corresponding to each port will turn on green to indicate that the connection is valid.

Connecting to SFP Fiber Optic Ports

The switch provides four slots for SFP-compliant fiber-optic transceivers.

Note that all 1000BASE fiber optic ports operate at 1 Gbps full duplex. All 100BASE fiber optic ports operate at 100 Mbps full duplex.

Maximum Gigabit Ethernet Fiber Cable Lengths			
Cable Type	Fiber Bandwidth	Maximum Cable Length	Connector
1000BASE-SX			
62.5/125 micron multimode	160 MHz/km	2-220 m (7-722 ft)	LC
	200 MHz/km	2-275 m (7-902 ft)	LC
50/125 micron multimode	400 MHz/km	2-500 m (7-1641 ft)	LC
	500 MHz/km	2-550 m (7-1805 ft)	LC
1000BASE-LX			
9/125 micron single-mode	N/A	2 m-10 km (7 ft-6.2 miles)	LC
1000BASE-LH			
9-125 micro single-mode	N/A	2 m-80 km (7 ft-50 miles)	LC
1000BASE-BX10			
9-125 micro single-mode	simplex fiber	2 m-10 km (7 ft-6.2 miles)	LC
1000BASE-BX20			
9-125 micro single-mode	simplex fiber	2 m-20 km (7 ft-12.4 miles)	LC
100BASE-BX20			
9-125 micro single-mode	simplex fiber	2 m-20 km (7 ft-12.4 miles)	LC

NOTE: The length of fiber optic cable for a single switched link should not exceed the relevant standards specified in this section. However, power budget constraints should also be considered when calculating the maximum fiber optic cable length for a particular link.

NOTE: Maximum distances may vary for different SFP vendors.

Follow these steps to connect cables to SFP transceiver ports:

WARNING: This switch uses lasers to transmit signals over fiber optic cable. The lasers are compliant with the requirements of a Class 1 Laser Product and are inherently eye safe in normal operation. However, you should never look directly at a transmit port when it is powered on.

WARNING: When selecting a fiber SFP device, considering safety, please make sure that it can function at a temperature that is not less than the recommended maximum operational temperature of the product. You must also use an approved Laser Class 1 SFP transceiver.

1. Remove and keep the fiber port's rubber plug. When not connected to a fiber cable, the rubber plug should be replaced to protect the optics.
2. Check that the fiber terminators are clean. You can clean the cable plugs by wiping them gently with a clean tissue or cotton ball moistened with a small amount of ethanol. Dirty fiber terminators on fiber optic cables will impair the quality of the light transmitted through the cable and lead to degraded performance on the port.
3. Connect one end of the cable to the SFP port on the switch and the other end to the SFP port on the other device. Since SFP connectors are keyed, the cable can be attached in only one orientation.



FIG. 15 Making a Connection to an SFP Port

4. As a connection is made, check the Link LED on the switch to be sure that the connection is valid.

NOTE: Be sure to secure cables properly and route them away from the switch without exceeding the minimum bending radius for fiber cables (typically a few inches). Use cable ties to bundle cables together and secure coiled loops of excess cable. Do not let cables hang free supporting their own weight or pull in any way that puts stress on the connectors.

Switch Management

The switches include a management agent that allows you to configure or monitor the switch using its embedded management software. To manage the switch, you can make a direct connection to the console port (out-of-band), or you can manage it through a network connection (in-band) using Telnet, Secure Shell (SSH), a web browser, or SNMP-based network management software. For a detailed description of the switch's software features, refer to the Web Management chapters later in this guide and *CLI Reference Guide*.

Understanding the System Status LEDs

The switch includes a display panel of key system LED indicators.



FIG. 16 System Status LEDs

The front panel LEDs are described in the following table:

System Status LEDs		
LED	Condition	Status
Power/PoE Max	On Green	Internal power operating normally.
	On Amber	The PoE device power draw on the switch has reached the system limitation.
	Off	No AC power is connected or the internal power supply has failed.
Diag (Diagnostic)	On Green	The system diagnostic test has been completed successfully.
	Flashing Green	System diagnostic in progress.
	Off	System boot up failed.

Connecting to the Console Port

The RJ-45 Console port on the front panel of the switch is used to connect a console device to the switch for out-of-band console configuration. The console device can be a PC or workstation running a VT-100 terminal emulator, or a VT-100 terminal. A console cable is supplied with the switch for connecting to a PC's RS-232 serial DB-9 DTE (COM) port.

NOTE: To connect to notebooks or other PCs that do not have a DB-9 COM port, use a USB-to-male DB-9 adapter cable (not included with the switch).




FIG. 17 Console Port

The following table describes the pin assignments used in the console cable:

Console Cable Wiring		
Switch's RJ-45 Console Port	Null Modem	PC's 9-Pin DTE Port
6 RXD (receive data)	<-----	3 TXD (transmit data)
3 TXD (transmit data)	----->	2 RXD (receive data)
4, 5 SGND (signal ground)	-----	5 SGND (signal ground)

No other pins are used.

The serial port's default settings are as follows:

- Default Baud rate: 115200 bps
- Character Size: 8 Characters
- Parity: None
- Stop bit: One
- Data bits: 8
- Flow control: None



Console cable to Console port → Connect to DB-9 COM port on PC

FIG. 18 Console Port Connection

Follow these steps to connect to the Console port:

1. Connect one end of the included RJ-45 to DB-9 serial cable to a DB-9 COM port connector on a management PC.
2. Plug in the RJ-45 end of the serial cable to the Console port on the switch.
3. Configure the PC's COM port required settings using VT-100 terminal emulator software (such as HyperTerminal) running on the management PC. The switch's default console port settings are:
115200 bps, 8 data bits, 1-stop bit and no parity
4. Log in to the command-line interface (CLI) using one of the default user login settings:
User - admin,
Password - admin
or
User - guest
Password - guest

The switch also offers a user-friendly web-based management interface for the configuration of all the unit's features.

You can make initial configuration changes by connecting a PC directly to one of the switch's RJ-45 ports. The switch has a default management IP address of 192.168.2.10 and a subnet mask of 255.255.255.0. You must set your PC IP address to be on the same subnet as the switch (that is, the PC and switch addresses must both start 192.168.2.x).

Log in to the web interface using the default settings:

Login Name - admin
Password - admin

Note that the guest default user login will only allow a user to view switch parameter data. For a detailed description of connecting to the console and using the switch's command line interface (CLI), refer to the *CLI Reference Guide*.

Resetting the Switch

The Reset button located on the front right side panel of the switch can be used to restart the device and set the configuration back to either the currently saved configuration or the factory default settings.

Resetting to the Saved Configuration File

Press the Reset button for less than 5 seconds to restart the system software using the current saved configuration file settings. Any unsaved changes in the currently running configuration will be lost and the only the saved settings in the startup configuration file will be used when the switch reboots.

Resetting to the Factory Default Settings

Press the Reset button for more than 5 seconds to restart the system software using the factory default settings. Any unsaved changes in the currently running configuration will be lost. The saved startup configuration file will still be available to select within the switch user interface, if needed.

CAUTION: Pressing the reset button will lose any unsaved changes in the running switch configuration.




FIG. 19 Reset Button

Using the Web Interface

Overview

The NXA-ENET8-POE+ 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 9, Mozilla Firefox 39, or Google Chrome 44, or more recent versions).

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

Connecting to the Web Interface

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

1. Configure the switch with a valid IP address, subnet mask, and default gateway using an out-of-band serial connection, BOOTP or DHCP protocol. (See the *Initial Switch Configuration* section 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 on-board configuration program. (See the *Configuring User Accounts* section on page 144.)
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 the *Configuring Interface Settings for STA* section on page 108 for more information.

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

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

Navigating the Web Browser Interface

To access the web-browser interface you must first enter a user name and password. The administrator has Read/Write access to all configuration parameters and statistics. The default user name and password for the administrator is "admin." The administrator has full access privileges to configure any parameters in the web interface. The default user name and password for guest access is "guest." The guest only has read access for most configuration parameters. Refer to the *Configuring User Accounts* section on page 144 for more details.

Dashboard

When your web browser connects with the switch's web agent, the Dashboard is displayed as shown below. The Dashboard displays the main menu on the left side of the screen. Switch Information, CPU Utilization, Switch Events, Memory Utilization, Recent 5 Event Information, Port Utilization, Dynamic Address Count, and LLDP Remote Device Port List are displayed on the right side. The main menu links are used to navigate to other menus, and display configuration parameters and statistics.

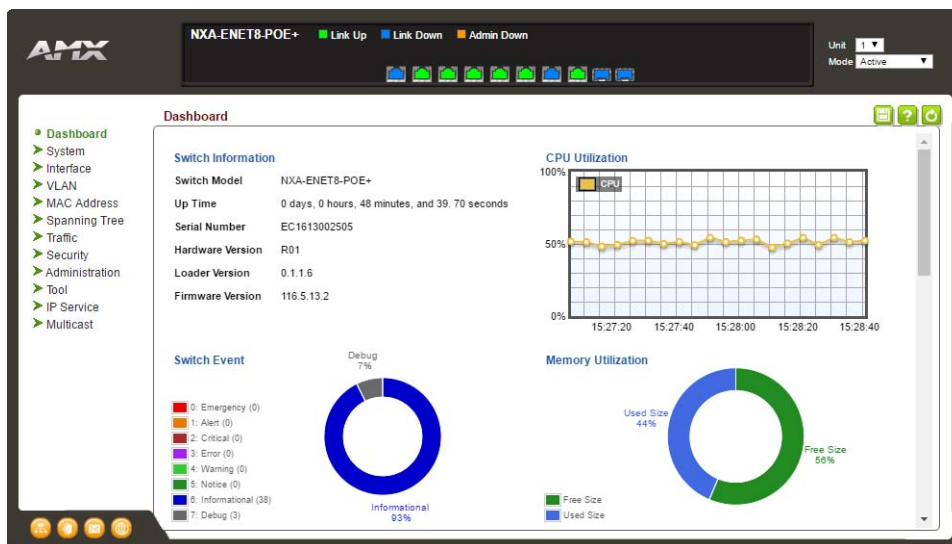


FIG. 20 Dashboard

NOTE: You can open a connection to the vendor's web site by clicking the AMX logo.

Home Page

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

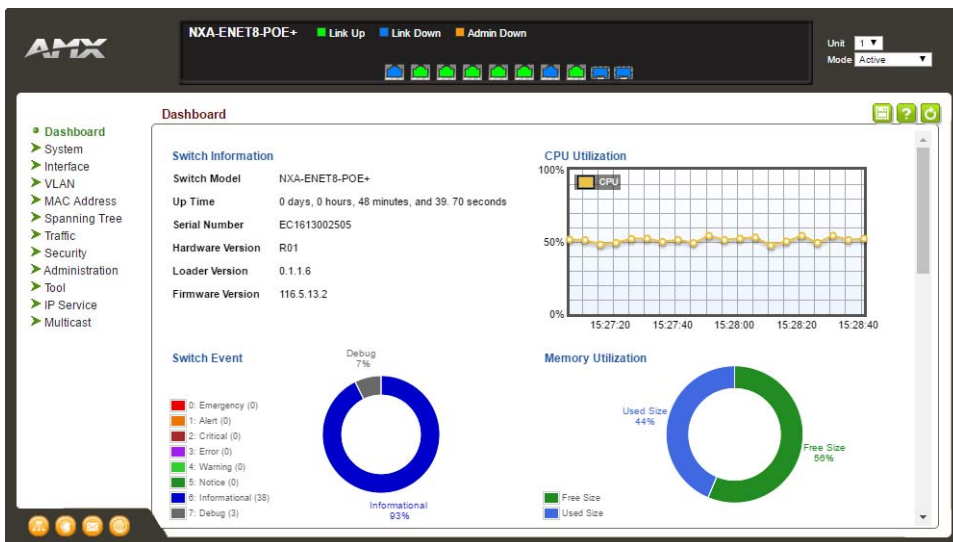









FIG. 21 Home Page

NOTE: You can open a connection to the vendor's web site by clicking the AMX 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.

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

Panel Display

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




FIG. 22 Front Panel Indicators

Basic Management Tasks

Overview

This chapter describes the following topics:

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

System Information

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

The following table lists the options on this page:

System - General Options	
System Description	A brief description of the device type.
System Object ID	MID II object ID for the switch's network management subsystem.
System Up Time	Length of time the management agent has been up.
System Name	Name assigned to the switch system.
System Location	Specifies the system location.
System Contact	Administrator responsible for the system.

Perform the following steps to configure general system information:

1. Click **System > General**. The System Information page opens (FIG. 23).

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

- System Description:** NXA-ENET8-POE+
- System Object ID:** 1.3.6.1.4.1.47645.43.103
- System Up Time:** 0 days, 1 hours, 7 minutes, and 10.60 seconds
- System Name:** [Empty input field]
- System Location:** [Empty input field]
- System Contact:** [Empty input field]
- Buttons:** Apply, Revert

FIG. 23 System Information

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

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.

The following table lists the options on this page. The options on this page are view-only.

System - Switch Options	
Main Board Information	
Serial Number	Displays the serial number of the switch.
Number of Ports	Displays the number of built-in ports.
Hardware Version	Displays the hardware version of the main board.
Main Power Status	Displays the status of the internal power supply.
Management Software Information	
Role	Displays whether the switch is operating as Master or Slave.
Loader Version	Displays the version number of the loader code.
Operation Code Version	Displays the version number of the runtime code.

To view the hardware and software version information, click **System**, then **Switch**.

System > Switch	
Main Board Information	
Serial Number	EC1613002505
Number of Ports	10
Hardware Version	R01
Main Power Status	Up
Management Software Information	
Role	Master
Loader Version	0.1.1.6
Operation Code Version	116.5.13.2

FIG. 24 General Switch Information

System Capabilities

Use the System > Capability page to configure jumbo frames or bridge extension settings.

System - Capability Options	
General Capability	
Jumbo Frame	Configures support for jumbo frames. The default setting is disabled.
Bridge Extension	
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 193.)
Static Entry Individual Port	This switch allows static filtering for unicast and multicast addresses. (Refer to "Setting Static Addresses" on page 157.)
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 139.)
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.

Configuring Support for Jumbo Frames

Use the System > Capability page to configure support for layer 2 jumbo frames. The switch provides more efficient throughput for large sequential data transfers by supporting jumbo frames up to 10240 bytes for Gigabit Ethernet and 10 Gigabit Ethernet ports or trunks. Compared to standard Ethernet frames that run only up to 1.5 KB, using jumbo frames significantly reduces the per-packet overhead required to process protocol encapsulation fields.

Usage Guidelines

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

The Capability page displays the Jumbo Frame option. By default, this option is disabled.

Perform these steps to configure support for jumbo frames:

1. Click **System > Capability**.
2. Enable or disable support for jumbo frames.
3. Click **Apply**.




FIG. 25 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.

The following table displays the options on this page:

System - Capability Options	
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 the <i>Class of Service</i> section on page 119.)
Static Entry Individual Port	This switch allows static filtering for unicast and multicast addresses. (Refer to the <i>Setting Static Addresses</i> section on page 97.)
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 the <i>VLAN Configuration</i> section on page 87.)
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.

To view Bridge Extension information, click **System**, then **Capability**.

System > Capability

General Capability

Jumbo Frame Enabled

Bridge Extension

Extended Multicast Filtering Services No

Traffic Classes Enabled

Static Entry Individual Port Yes

VLAN Version Number 2

VLAN Learning F/L

Local VLAN Capable No

Configurable PVID Tagging Yes

Max Supported VLAN Numbers 4094

Max Supported VLAN ID 4094

Apply Revert

FIG. 26 Displaying Bridge Extension Configuration

System Files

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

Copying Files via FTP/TFTP or HTTP

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

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

Command Usage

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

The following table displays the options on this page:

System - File (Copy) Options	
Copy Type	The firmware copy operation include the following options: <ul style="list-style-type: none"> • HTTP Upload - Copies a file from a management station to the switch. • HTTP Download - Copies a file from the switch to a management station • TFTP Upload - Copies a file from a TFTP server to the switch. • TFTP Download - Copies a file from the switch to a TFTP server. • FTP Upload - Copies a file from an FTP server to the switch. • FTP Download - Copies a file from the switch to an FTP server.
FTP/TFTP Server IP Address	The IP address of an FTP/TFTP server.
User Name	The user name for FTP server access.
Password	The password for FTP server access.
File Type	Specify Operation Code to copy firmware or Config File to copy configuration settings.
File Name	The file name should not contain slashes (\ or /), the leading letter of the file name should not be a period (.), and the maximum length for file names is 32 characters for files on the switch or 127 characters for files on the server. (Valid characters: A-Z, a-z, 0-9, ".", "-", "_").

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

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

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

Perform these steps to copy firmware files:

1. Click **System > File**.
2. Select **Copy** from the Action list.
3. Select **FTP Upload**, **HTTP Upload**, or **TFTP Upload** as the file transfer method.
4. If FTP or TFTP Upload is used, enter the IP address of the file server.
5. If FTP Upload is used, enter the user name and password for your account on the FTP server.
6. Set the file type to **Operation Code**.

7. Enter the name of the file to download.
8. Select a file on the switch to overwrite or specify a new file name.
9. Click **Apply**.

FIG. 27 Copy Firmware

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

Saving the Running Configuration to a Local File

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

The following table displays the options on this page:

System - File (Copy) Options	
Copy Type	The firmware copy operation include the following option: <ul style="list-style-type: none"> • Running-Config - Copies the current configuration settings to a local file on the switch.
Destination File Name	Copy to the currently designated startup file, or to a new file. The file name should not contain slashes (\ or /), the leading letter of the file name should not be a period (.), and the maximum length for file names is 32 characters. (Valid characters: A-Z, a-z, 0-9, ".", "-", "_").

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

Perform these steps to save the running configuration file:

1. Click **System > 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. Click **Apply**.

FIG. 28 Saving the Running Configuration

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. Perform these steps to set a file to use for system initialization:

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

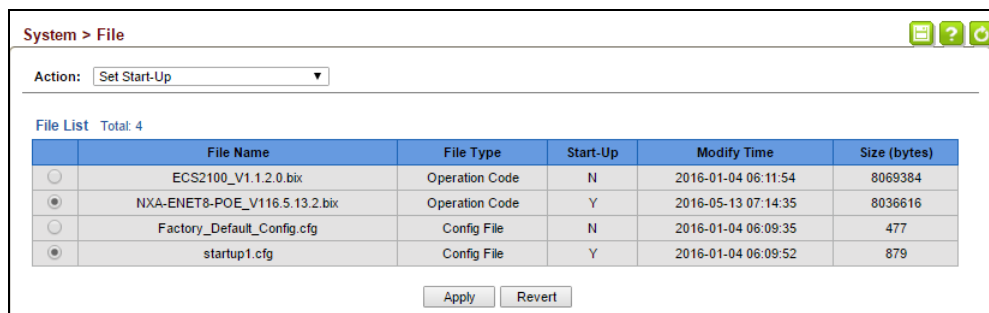


FIG. 29 Setting Start-Up Files

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.

Perform these steps to show the system files:

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

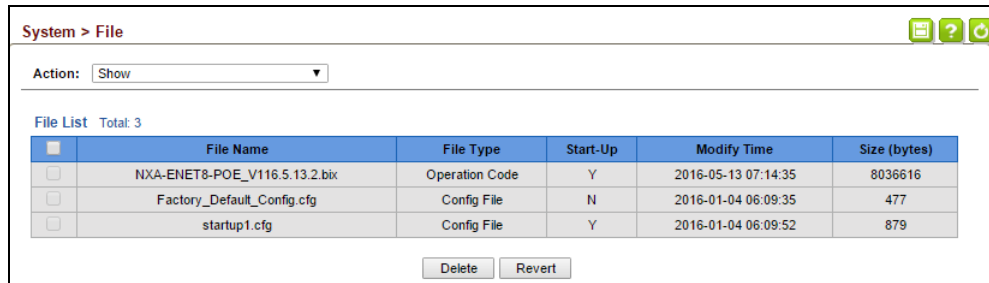


FIG. 30 Displaying System Files

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 ECS2100-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 ECS2100-Series.BIX from the server even though ECS2100-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, ecs2100-series.bix and ECS2100-Series.bix are considered to be unique files. Thus, if the upgrade file is stored as ECS2100-Series.bix (or even EcS2100-Series.bix) on a case-sensitive server, then the switch (requesting ecs2100-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.

The following parameters are displayed:

System - File Options	
Automatic Opcode Upgrade	Enables the switch to search for an upgraded operation code file during the switch bootup process. By default, this option is disabled.
Automatic Upgrade Location URL	<p>Defines where the switch should search for the operation code upgrade file. The last character of this URL must be a forward slash ("/"). The .bix filename must not be included since it is automatically appended by the switch. The available options are FTP and TFTP.</p> <p>The following syntax must be observed:</p> <pre>tftp://host[/filedir]/</pre> <ul style="list-style-type: none"> • 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. <pre>ftp://[username[:password@]]host[/filedir]/</pre> <ul style="list-style-type: none"> • 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. <p>Examples:</p> <p>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:</p> <ul style="list-style-type: none"> • tftp://192.168.0.1/ <p>The image file is in the TFTP root directory.</p> <ul style="list-style-type: none"> • tftp://192.168.0.1/switch-opcode/ <p>The image file is in the "switch-opcode" directory, relative to the TFTP root.</p> <ul style="list-style-type: none"> • tftp://192.168.0.1/switches/opcode/ <p>The image file is in the "opcode" directory, which is within the "switches" parent directory, relative to the TFTP root.</p> <p>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:</p> <ul style="list-style-type: none"> • ftp://192.168.0.1/ <p>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.</p> <ul style="list-style-type: none"> • ftp://switches:upgrade@192.168.0.1/ <p>The user name is "switches" and the password is "upgrade". The image file is in the FTP root.</p> <ul style="list-style-type: none"> • ftp://switches:upgrade@192.168.0.1/switches/opcode/ <p>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.</p>

Perform these steps to configure automatic code upgrade:

1. Click **System > 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**.

FIG. 31 Configuring Automatic Code Upgrade

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

```
Automatic Upgrade is looking for a new image
New image detected: current version 1.2.1.3; new version 1.2.1.6
Image upgrade in progress
The switch will restart after upgrade succeeds
Downloading new image
```

```
Flash programming started
Flash programming completed
The switch will now restart
```

Setting the System Clock

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

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

Setting the Time Manually

Use the System > Time (Configure General - Manual) page to set the system time on the switch manually without using SNTP. The following table displays the options on this page:

System - Time Options	
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)

Perform these steps to manually set the system clock:

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

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

FIG. 32 Manually Setting the System Clock

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.

The following table displays the options on this page:

System - Time Options	
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)

Perform these steps to set the polling interval for SNTP:

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

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

FIG. 33 Setting the Polling Interval for SNTP

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.

The following table displays the options on this page:

System - Time Options	
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)

Perform these steps to set the clock maintenance type to NTP:

1. Click **System > 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**.

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 '2014-5-30 9:59:20'. The 'Maintain Type' is set to 'NTP' via a dropdown menu. Under the 'NTP Configuration' section, the 'Authentication Status' is currently 'Disabled' (indicated by an unchecked checkbox). The 'Polling Interval' is set to '1024 sec'. At the bottom right of the configuration area, there are 'Apply' and 'Revert' buttons.

FIG. 34 Configuring NTP

Configuring Time Servers

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

Specifying SNTP Time Servers

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

The following table displays the options on this page:

System - Time Options	
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.

Perform these steps to set the SNTP time servers:

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

FIG. 35 Specifying SNTP Time Server

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.

The following table displays the options on this page:

System - Time Options	
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.
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)

Perform these steps to add an NTP time server to the server list:

1. Click **System > 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**.

FIG. 36 Adding an NTP Time Server

Perform these steps to show the list of configured NTP time servers:

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

FIG. 37 Showing the NTP Time Server List

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. The following table displays the options on this page:

System - Time Options	
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.

Perform these steps to add an entry to NTP authentication key list:

1. Click **System > 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**.

FIG. 38 Adding an NTP Authentication Key

Perform these steps to show the list of configured NTP authentication keys:

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

NTP Authentication Key List Total: 1	
Authentication Key	Key Context
3	BJ0774Q8899747D10867F125505J62770084706278G1357878N8475052113Q89137L8

FIG. 39 Showing the NTP Authentication Key List

Setting the Time Zone

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

The following table displays the options on this page:

System - Time Options	
Predefined Configuration	A drop-down box provides access to the 80 predefined time zone configurations. Each choice indicates its offset from UTC and lists at least one major city or location covered by the time zone.
User-defined Configuration	<p>Allows the user to define all parameters of the local time zone.</p> <ul style="list-style-type: none"> • Direction - Configures the time zone to be before (east of) or after (west of) UTC. • Name - Assigns a name to the time zone. (Range: 1-30 characters) • Hours (0-13) - The number of hours before or after UTC. The maximum value before UTC is 12. The maximum value after UTC is 13. • Minutes (0-59) - The number of minutes before/after UTC.

Perform these steps to set your local time zone:

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

The screenshot shows the 'System > Time' configuration page. At the top, there is a breadcrumb 'System > Time' and a 'Step:' dropdown menu set to '3. Configure Time Zone'. Below this, there are two radio button options: 'Predefined Configuration' (which is unselected) and 'User Defined Configuration' (which is selected). Under 'User Defined Configuration', there are four input fields: 'Direction' (a dropdown menu set to 'After UTC'), 'Name' (a text box containing 'UTC'), 'Hours (0-13)' (a text box containing '0'), and 'Minutes (0-59)' (a text box containing '0'). Below these fields is a note: 'Note: The maximum value before UTC is 12:00. The maximum value after UTC is 13:00.' At the bottom right of the form are two buttons: 'Apply' and 'Revert'.

FIG. 40 Setting the Time Zone

Configuring Summer Time

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

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

The following table displays the options on this page:

System - Time Options	
General Configuration	
Summer Time in Effect	Shows if the system time has been adjusted.
Status	Shows if summer time is set to take effect during the specified period.
Name	Name of the time zone while summer time is in effect, usually an acronym. (Range: 1-30 characters)
Mode	<p>Selects one of the following configuration modes. (The Mode option can only be managed when the Summer Time Status option has been set to enabled for the switch.)</p> <ul style="list-style-type: none"> • Predefined Mode - Configures the summer time status and settings for the switch using predefined configurations for several major regions of the world. To specify the time corresponding to your local time when summer time is in effect, select the predefined summer-time zone appropriate for your location. • Date Mode - Sets the start, end, and offset times of summer time for the switch on a one-time basis. This mode sets the summer-time zone relative to the currently configured time zone. To specify a time corresponding to your local time when summer time is in effect, you must indicate the number of minutes your summer- time zone deviates from your regular time zone. <ul style="list-style-type: none"> Offset - Summer-time offset from the regular time zone, in minutes. (Range: 1-120 minutes) From - Start time for summer-time offset. To - End time for summer-time offset. • Recurring Mode - Sets the start, end, and offset times of summer time for the switch on a recurring basis. This mode sets the summer-time zone relative to the currently configured time zone. To specify a time corresponding to your local time when summer time is in effect, you must indicate the number of minutes your summer- time zone deviates from your regular time zone. <ul style="list-style-type: none"> Offset - Summer-time offset from the regular time zone, in minutes. (Range: 1-120 minutes) From - Start time for summer-time offset. To - End time for summer-time offset.

Perform these steps to specify summer time settings:

1. Click **System > Time**.
2. Select **Configure Summer Time** from the Step list.
3. Select one of the configuration modes, configure the relevant attributes, enable summer time status.
4. Click **Apply**.

The screenshot shows the 'System > Time' configuration page. At the top, it says 'Step: 4. Configure Summer Time'. Below this, there are several configuration options:

- Summer Time in Effect:** Set to 'No'.
- Status:** A checkbox labeled 'Enabled' is checked.
- Name:** An empty text input field.
- Mode:** A dropdown menu set to 'Predefined'.
- Predefined Mode Configuration:** A section with a dropdown menu for 'Daylight Savings' set to 'Australia'.

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

FIG. 41 Configuring Summer Time

Configuring the Console Port

Use the System > Console menu to configure connection parameters for the switch's console port. You can access the on-board 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. The following table displays the options on this page:

System - Console Options	
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 login attempts. When the login 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 login attempt. (Range: 1-120; Default: 3 attempts)
Silent Time	Sets the amount of time the management console is inaccessible after the number of unsuccessful login attempts has been exceeded. (Range: 1-65535 seconds; Default: Disabled)
Data Bits	Sets the number of data bits per character that are interpreted and generated by the console port. If parity is being generated, specify 7 data bits per character. If no parity is required, specify 8 data bits per character. (Default: 8 bits)
Stop Bits	Sets the number of the stop bits transmitted per byte. (Range: 1-2; Default: 1 stop bit)
Parity	Defines the generation of a parity bit. Communication protocols provided by some terminals can require a specific parity bit setting. Specify Even, Odd, or None. (Default: None)
Speed	Sets the terminal line's baud rate for transmit (to terminal) and receive (from terminal). Set the speed to match the baud rate of the device connected to the serial port. (Range: 9600, 19200, 38400, 57600, or 115200 baud; Default: 115200 baud)

NOTE: The password for the console connection can only be configured through the CLI (see the "password" command in the CLI Reference Guide).

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

Perform these steps to configure parameters for the console port:

1. Click **System > Console**.
2. Specify the connection parameters as required.
3. Click **Apply**.

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

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

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

FIG. 42 Console Port Settings

Configuring Telnet Settings

Use the System > Telnet menu to configure parameters for accessing the CLI over a Telnet connection. You can access the on-board 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.

The following table displays the options on this page:

System - Time Options	
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 login attempts. When the login 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 login attempt. (Range: 1-120; Default: 3 attempts)
Silent Time	Sets the amount of time the management interface is inaccessible after the number of unsuccessful login attempts has been exceeded. (Range: 1-65535 seconds; Default: Disabled)

NOTE: The password for the Telnet connection can only be configured through the CLI (see the "password" command in the CLI Reference Guide).

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

Perform these steps to configure parameters for the console port:

1. Click **System > Telnet**.
2. Click the **Enabled** check box to enable Telnet on the switch.
3. Specify the connection parameters as required.
4. Click **Apply**.

The screenshot shows the 'System > Telnet' configuration page. The settings are as follows:

- Telnet Status: Enabled
- TCP Port (1-65535):
- Max Sessions (0-8):
- Login Timeout (10-300): sec
- Exec Timeout (60-65535): sec
- Password Threshold (1-120):
- Silent Time (1-65535): sec

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

FIG. 43 Telnet Connection Settings

Displaying CPU Utilization

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

The following table displays the options on this page:

System - CPU Utilization Options	
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

Perform the following steps to display CPU utilization:

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




FIG. 44 Displaying CPU Utilization

Configuring CPU Guard

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

The following table displays the options on this page:

System - CPU Guard Options	
CPU Guard Status	Enables CPU Guard. (Default: Disabled)
High Watermark	If the percentage of CPU usage time is higher than the high-watermark, the switch stops packet flow to the CPU (allowing it to catch up with packets already in the buffer) until usage time falls below the low watermark. (Range: 40-100%; Default: 90%)
Low Watermark	If packet flow has been stopped after exceeding the high watermark, normal flow will be restored after usage falls beneath the low watermark. (Range: 40-100%; Default: 70%)
Maximum Threshold	If the number of packets being processed by the CPU is higher than the maximum threshold, the switch stops packet flow to the CPU (allowing it to catch up with packets already in the buffer) until the number of packets being processed falls below the minimum threshold. (Range: 50-500 pps; Default: 500 pps)
Minimum Threshold	If packet flow has been stopped after exceeding the maximum threshold, normal flow will be restored after usage falls beneath the minimum threshold. (Range: 50-500 pps; Default: 50 pps)
Trap Status	If enabled, an alarm message will be generated when utilization exceeds the high watermark or exceeds the maximum threshold. (Default: Disabled) Once the high watermark is exceeded, utilization must drop beneath the low watermark before the alarm is terminated, and then exceed the high watermark again before another alarm is triggered. Once the maximum threshold is exceeded, utilization must drop beneath the minimum threshold before the alarm is terminated, and then exceed the maximum threshold again before another alarm is triggered.
Current Threshold	Shows the configured threshold in packets per second.

Perform these steps to configure CPU Guard:

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

System > CPU Guard

CPU Guard Status Enabled

High Watermark (40-100) %

Low Watermark (40-100) %

Maximum Threshold (50-500) packets/sec

Minimum Threshold (50-500) packets/sec

Trap Status Enabled

Current Threshold 500 packets/sec

FIG. 45 Configuring CPU Guard

Displaying Memory Utilization

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

The following table displays the options on this page:

System - Memory Status Options	
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.

To display memory utilization, click **System**, then **Memory Status**.

System > Memory Status

Memory Status

Free Size	45,416,448 bytes	16%
Used Size	223,019,008 bytes	84%
Total	268,435,456 bytes	

FIG. 46 Displaying Memory Utilization

Resetting the System

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

Command Usage

- This command resets the entire system.
- When the system is restarted, it will always run the Power-On Self-Test. It will also retain all configuration information stored in non-volatile memory. (See the *Saving the Running Configuration to a Local File* section on page 43 for more information.)

The following table displays the options on this page:

System - Time Options	
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	<p>Restarts the switch immediately or at the specified time(s).</p> <ul style="list-style-type: none"> • 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.) <ul style="list-style-type: none"> 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. <ul style="list-style-type: none"> 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. <ul style="list-style-type: none"> Time: <ul style="list-style-type: none"> HH - The hour at which to reload. (Range: 00-23) MM - The minute at which to reload. (Range: 00-59) Period: <ul style="list-style-type: none"> 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)

Perform these steps to restart the switch:

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

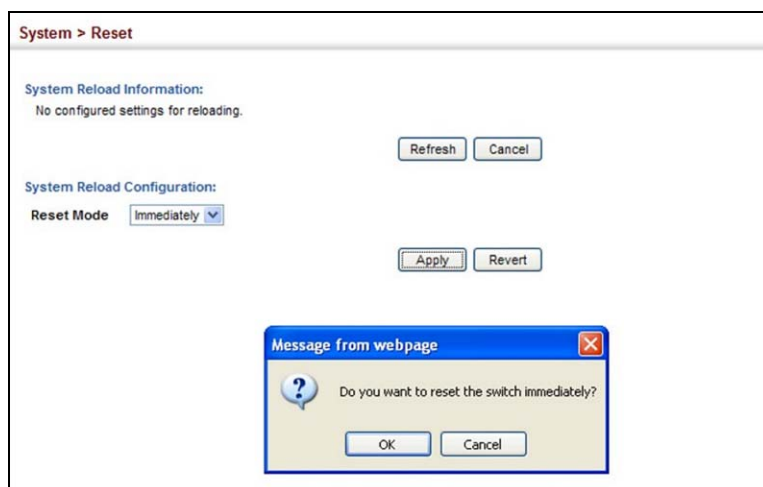


FIG. 47 Restarting the Switch (Immediately)

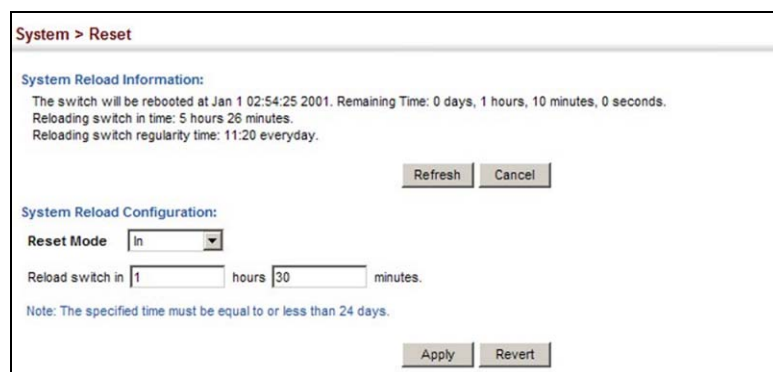


FIG. 48 Restarting the Switch (In)

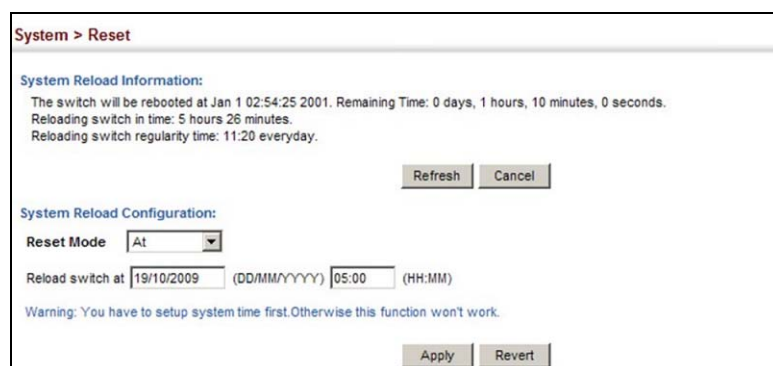


FIG. 49 Restarting the Switch (At)

System > Reset

System Reload Information:
No configured settings for reloading.

System Reload Configuration:

Reset Mode Regularly

Time 05:30 (HH:MM)

Period

Daily

Weekly

Monthly

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

FIG. 50 Restarting the Switch (Regularly)

Interface Configuration

Overview

This chapter describes the following topics:

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

Port Configuration

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

Configuring by Port List

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

Command Usage

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

NOTE: Auto-negotiation is not supported for 1000BASE SFP transceivers.

The following table lists the options on this page:

Interface - Port (General) Options	
Port	Port identifier. (Range: 1-10/26/28/52)
Type	Indicates the port type. (1000BASE-T or 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. (Default: Enabled)
Autonegotiation (Port Capabilities)	<p>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.</p> <ul style="list-style-type: none"> • 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 - Symmetric exchange of 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. <p>Default: Auto-negotiation enabled; Advertised capabilities for 100BASE-FX (SFP) - 100full 1000BASE-T - 10half, 10full, 100half, 100full, 1000full 1000BASE-SX/LX/ZX (SFP) - 1000full</p>
Speed/Duplex	Allows you to manually set the port speed and duplex mode. (i.e., with auto-negotiation disabled)

Interface - Port (General) Options	
Flow Control	Allows automatic or manual selection of flow control. (Default: Enabled)
Link Up Link Down	Issues a notification message whenever a port link is established or broken. (Default: Disabled)

Perform these steps 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**.

FIG. 51 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.

Except for the trap command, refer to the *Configuring by Port List* section on page 61 for more information on command usage and a description of the options on the page.

Perform these steps to configure port connection parameters:

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

FIG. 52 Configuring Connections by Port Range

Displaying Connection Status

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

The following table lists the options on this page:

Interface - General (Show Information) Options	
Port	Port identifier. (Range: 1-10/26/28/52)
Type	Indicates the port type. (1000BASE-T 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
Shutdown Reason	Shows the reason this interface has been shut down if applicable. Some of the reasons for shutting down an interface include being administratively disabled, or exceeding traffic boundary limits set by auto traffic control.
Autonegotiation (Port Capabilities)	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
Link Up Link Down	Shows if a notification message will be sent whenever a port link is established or broken. (Default: Enabled)

Perform these steps to display port connection parameters:

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

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

FIG. 53 Displaying Port Information

Showing Port or Trunk Statistics

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

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

The following table lists the options on this page:

Port Statistics	
Interface Statistics	
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.

Port Statistics	
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.
Etherlike Statistics	
Single Collision Frames	The number of successfully transmitted frames for which transmission is inhibited by exactly one collision.
Multiple Collision Frames	A count of successfully transmitted frames for which transmission is inhibited by more than one collision.
Late Collisions	The number of times that a collision is detected later than 512 bit-times into the transmission of a packet.
Excessive Collisions	A count of frames for which transmission on a particular interface fails due to excessive collisions. This counter does not increment when the interface is operating in full-duplex mode.
Deferred Transmissions	A count of frames for which the first transmission attempt on a particular interface is delayed because the medium was busy.
Frames Too Long	A count of frames received on a particular interface that exceed the maximum permitted frame size.
Alignment Errors	The number of alignment errors (mis-synchronized 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.
RMON Statistics	
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).

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

Perform these steps to show a list of port statistics:

1. Click **Interface > Port > Statistics**.
2. Select the statistics mode to display (Interface, Etherlike, RMON, or Utilization).
3. Select a port from the drop-down list.
4. Use the **Refresh** button to update the screen.

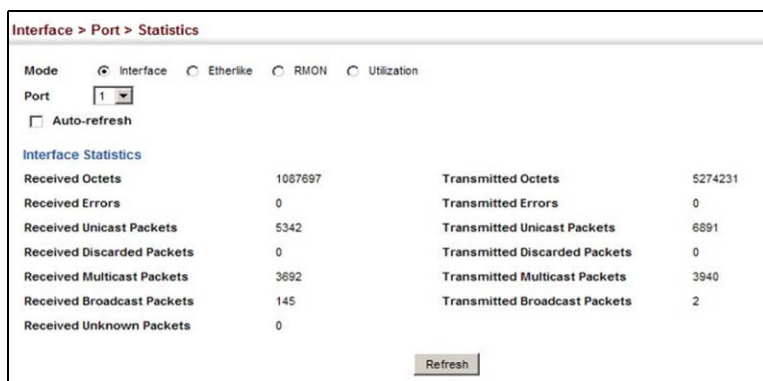


FIG. 54 Showing Port Statistics (Table)

Perform these steps 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.




FIG. 55 Showing Port Statistics (Chart)

Displaying Statistical History

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

Command Usage

- For a description of the statistics displayed on these pages, see the *Showing Port or Trunk Statistics* section on page 63.
- To configure statistical history sampling, use the *Displaying Statistical History* section on page 66.

The following table lists the options on this page:

Interface - History Options	
Add	
Port	Port number (Range: 1-10/26/28/52)
History Name	Name of sample interval (Range: 1-32 characters)
Interval	The interval for sampling statistics (Range: 1-86400 minutes)
Requested Buckets	The number of samples to take (Range: 1-96)
Show	
Port	Port number (Range: 1-10/28)
History Name	Name of sample interval (Default settings: 15min, 1day)
Interval	The interval for sampling statistics
Requested Buckets	The number of samples to take
Show Details	

Interface - History Options	
Mode	<ul style="list-style-type: none"> • Status - Shows the sample parameters. • Current Entry - Shows current statistics for the specified port and named sample. • Input Previous Entries - Shows statistical history for ingress traffic. • Output Previous Entries - Shows statistical history for egress traffic.
Port	Port number (Range: 1-10/28)
Name	Name of sample interval

Perform these steps to configure a periodic sample of statistics:

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

FIG. 56 Configuring a History Sample

Perform these steps to show the configured entries for a history sample:

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

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

FIG. 57 Showing Entries for History Sampling

Perform these steps to show the configured parameters for a sampling entry:

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

The screenshot shows the 'Interface > Port > History' page. At the top, there is an 'Action:' dropdown menu set to 'Show Details'. Below this, there are four radio buttons for 'Mode': 'Status' (selected), 'Current Entry', 'Input Previous Entries', and 'Output Previous Entries'. There are two dropdown menus: 'Port' set to '1' and 'Name' set to '15min'. Under the heading 'History Status', there is a table with the following data:

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

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

FIG. 58 Showing Status of Statistical History Sample

Perform these steps to show statistics for the current interval of a sample entry:

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

The screenshot shows the 'Interface > Port > History' page. At the top, there is an 'Action:' dropdown menu set to 'Show Details'. Below this, there are four radio buttons for 'Mode': 'Status', 'Current Entry' (selected), 'Input Previous Entries', and 'Output Previous Entries'. There are two dropdown menus: 'Port' set to '1' and 'Name' set to '15min'. Under the heading 'Current Entry', there is a table with the following data:

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

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

FIG. 59 Showing Current Statistics for a History Sample

Perform these steps to show ingress or egress traffic statistics for a sample entry:

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

Displaying Transceiver Data

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

The following table lists the options on this page:

Interface - Transceiver Options	
Port	Port number. (Range: 9-10/23-26/25-28/49-52)
General	Information on connector type and vendor-related parameters.
DDM Information	Information on temperature, supply voltage, laser bias current, laser power, and received optical power. The switch can display diagnostic information for SFP modules which support the SFF-8472 Specification for Diagnostic Monitoring Interface for Optical Transceivers. This information allows administrators to remotely diagnose problems with optical devices. This feature, referred to as Digital Diagnostic Monitoring (DDM) provides information on transceiver parameters.

Perform these steps to display identifying information and functional parameters for optical transceivers:

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

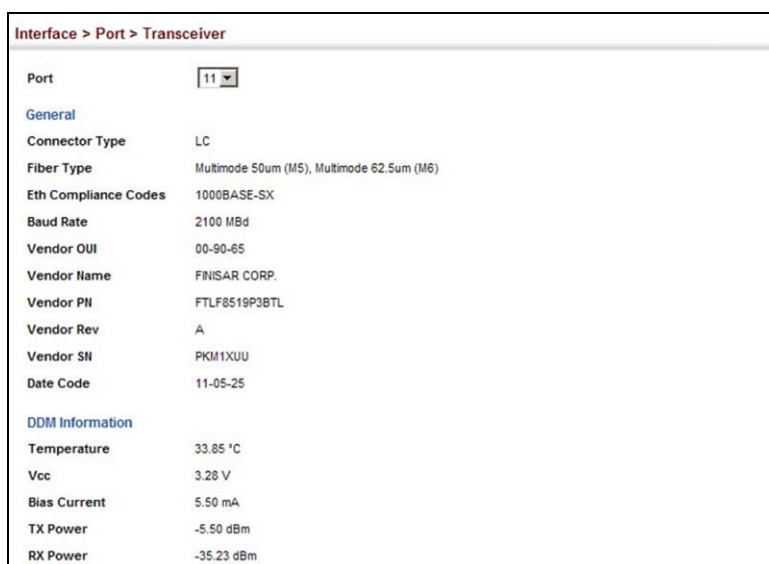


FIG. 60 Displays Transceiver Data

Configuring Transceiver Thresholds

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

The following table lists the options on this page:

Interface - Transceiver Options	
Port	Port number. (Range: 9-10/23-26/25-28/49-52)
General	Information on connector type and vendor-related parameters.
DDM Information	Information on temperature, supply voltage, laser bias current, laser power, and received optical power. The switch can display diagnostic information for SFP modules which support the SFF-8472 Specification for Diagnostic Monitoring Interface for Optical Transceivers. This information allows administrators to remotely diagnose problems with optical devices. This feature, referred to as Digital Diagnostic Monitoring (DDM) provides information on transceiver parameters.
Trap	Sends a trap when any of the transceiver's operation values falls outside of specified thresholds. (Default: Disabled)
Auto Mode	Uses default threshold settings obtained from the transceiver to determine when an alarm or trap message should be sent. (Default: Enabled)

Interface - Transceiver Options

DDM Thresholds

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

The following alarm and warning parameters are supported:

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

The configurable ranges are:

- Temperature: -128.00-128.00 °C
- Voltage: 0.00-6.55 Volts
- Current: 0.00-131.00 mA
- Power: -40.00-8.20 dBm

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

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

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

Perform these steps to configure threshold values for optical transceivers:

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

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

FIG. 61 Configuring Transceiver Thresholds

Trunk Configuration

This section describes how to configure static and dynamic trunks.

You can create multiple links between devices that work as one virtual, aggregate link. A port trunk offers a dramatic increase in bandwidth for network segments where bottlenecks exist, as well as providing a fault-tolerant link between two devices. You can create up to 16 trunks at a time on the switch, or up to 32 across the stack.

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

Command Usage

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

- Finish configuring trunks before you connect the corresponding network cables between switches to avoid creating a loop.
- You can create up to 16 trunks on a switch or 32 trunks in the stack, with up to eight ports per trunk.
- The ports at both ends of a connection must be configured as trunk ports.
- When configuring static trunks on switches of different types, they must be compatible with the Cisco EtherChannel standard.
- The ports at both ends of a trunk must be configured in an identical manner, including communication mode (i.e., speed, duplex mode and flow control), VLAN assignments, and CoS settings.
- Any of the Gigabit ports on the front panel can be trunked together, including ports of different media types.
- All the ports in a trunk have to be treated as a whole when moved from/to, added or deleted from a VLAN.
- STP, VLAN, and IGMP settings can only be made for the entire trunk.

Configuring a Static Trunk

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




FIG. 62 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.

The following table lists the options on this page:

Interface - Static Trunk Options	
Trunk ID	Trunk identifier (Range: 1-8)
Member	The initial trunk member. Use the Add Member page to configure additional members.
Unit	Unit identifier (Range: 1)
Port	Port identifier (Range: 1-10/26/28/52)

Perform these steps 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**.

FIG. 63 Creating Static Trunks

Perform these steps 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**.

FIG. 64 Adding Static Trunks Members

Perform these steps 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 the *Configuring by Port List* section on page 61 for a description of the parameters.)
5. Click **Apply**.

Trunk	Type	Name	Admin	Autonegotiation	Speed Duplex	Flow Control	Link Up Down Trap
1	1000BASE-T		Enabled	<input checked="" type="checkbox"/> Enabled <input checked="" type="checkbox"/> 10h <input checked="" type="checkbox"/> 100h <input checked="" type="checkbox"/> 1000F <input checked="" type="checkbox"/> 10F <input type="checkbox"/> Sym	1000uB	Enabled	Enabled

FIG. 65 Configuring Connection Parameters for a Static Trunk

Perform these steps 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.

Interface > Trunk > Static									
Step: 2. Configure General		Action: Show Information							
Static Trunk List Total: 1									
Trunk	Type	Name	Admin	Oper Status	Shutdown Reason	Autonegotiation	Oper Speed Duplex	Oper Flow Control	Link Up Down Trap
1	1000BASE-T		Enabled	Down		Enabled	1000Full	None	Enabled

FIG. 66 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.

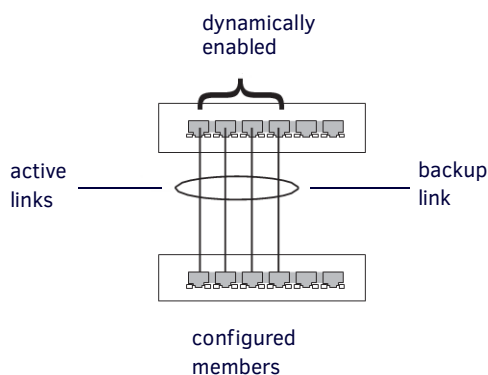


FIG. 67 Configuring Dynamic Trunks

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).

The following table lists the options on this page:

Interface - Static Trunk Options	
Configure Aggregator	
Admin Key	LACP administration key is used to identify a specific link aggregation group (LAG) during local LACP setup on the switch. (Range: 0-65535) If the port channel admin key is not set when a channel group is formed (i.e., it has the null value of 0), this key is set to the same value as the port admin key (see Configure Aggregation Port - Actor/Partner) used by the interfaces that joined the group. Note that when the LAG is no longer used, the port channel admin key is reset to 0. If the port channel admin key is set to a non-default value, the operational key is based upon LACP PDUs received from the partner, and the channel admin key is reset to the default value. The trunk identifier will also be changed by this process.

Interface - Static Trunk Options	
Timeout Mode	<p>The timeout to wait for the next LACP data unit (LACPDU):</p> <ul style="list-style-type: none"> • Long Timeout - Specifies a slow timeout of 90 seconds. (This is the default setting.) • Short Timeout - Specifies a fast timeout of 3 seconds. <p>The timeout is set in the LACP timeout bit of the Actor State field in transmitted LACPDU. When the partner switch receives an LACPDU set with a short timeout from the actor switch, the partner adjusts the transmit LACPDU interval to 1 second. When it receives an LACPDU set with a long timeout from the actor, it adjusts the transmit LACPDU interval to 30 seconds.</p> <p>If the actor does not receive an LACPDU from its partner before the configured timeout expires, the partner port information will be deleted from the LACP group.</p> <p>When a dynamic port-channel member leaves a port-channel, the default timeout value will be restored on that port.</p> <p>When a dynamic port-channel is torn down, the configured timeout value will be retained. When the dynamic port-channel is constructed again, that timeout value will be used.</p>
System Priority	LACP system priority is used to determine link aggregation group (LAG) membership, and to identify this device to other switches during LAG negotiations.
System MAC Address	The device MAC address assigned to each trunk.
Configure Aggregation Port - General	
Port	Port identifier (Range: 1-10/26/28/52)
LACP Status	Enables or disables LACP on a port.
Configure Aggregation Port - Actor/Partner	
Port	Port number (Range: 1-10/28)
Admin Key	<p>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)</p> <p>Once the remote side of a link has been established, LACP operational settings are already in use on that side. Configuring LACP settings for the partner only applies to its administrative state, not its operational state.</p> <p>NOTE: <i>Configuring the partner admin-key does not affect remote or local switch operation. The local switch just records the partner admin-key for user reference.</i></p> <p>By default, the actor's operational key is determined by port's link speed (1000f - 4, 100f - 3, 10f - 2), and copied to the admin key.</p>
System Priority	<p>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)</p> <p>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.</p>
Port Priority	<p>If a link goes down, LACP port priority is used to select a backup link. (Range: 0-65535; Default: 32768)</p> <ul style="list-style-type: none"> • 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.*

Perform these steps to configure the admin key for a dynamic trunk:

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

Trunk	Admin Key (0-65535)	Timeout Mode	System Priority	System MAC Address
1	0	Long Timeout	32768	00-ED-0C-00-00-FD
2	0	Long Timeout	32768	00-ED-0C-00-00-FD
3	0	Long Timeout	32768	00-ED-0C-00-00-FD
4	0	Long Timeout	32768	00-ED-0C-00-00-FD
5	0	Long Timeout	32768	00-ED-0C-00-00-FD

FIG. 68 Configuring the LACP Aggregator Admin Key

Perform these steps 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**.

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

FIG. 69 Enabling LACP on a Port

Perform these steps 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**.

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

FIG. 70 Configuring LACP Parameters on a Port

Perform these steps to show the active members of a dynamic trunk:

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




FIG. 71 Showing Members of a Dynamic Trunk

Perform these steps to configure connection parameters for a dynamic trunk:

1. Click **Interface > Trunk > Dynamic**.
2. Select **Configure Trunk** from the Step list.
3. Select **Configure** from the Action list.
4. Modify the required interface settings. (See the *Configuring by Port List* section on page 61 for a description of the interface settings.)
5. Click **Apply**.

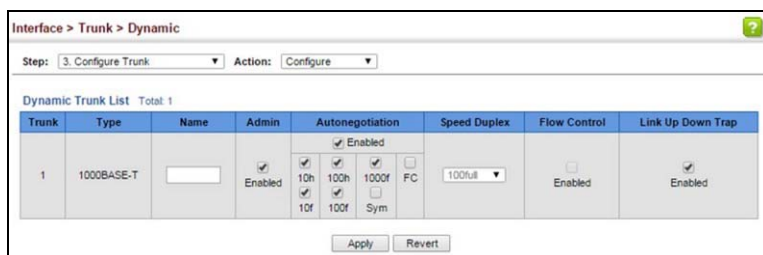


FIG. 72 Configuring Connection Settings for a Dynamic Trunk

Perform these steps to show connection parameters for a dynamic trunk:

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



FIG. 73 Showing Connection Parameters for 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.

The following table lists the options on this page:

LACP Port Counters	
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.

Perform these steps 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.

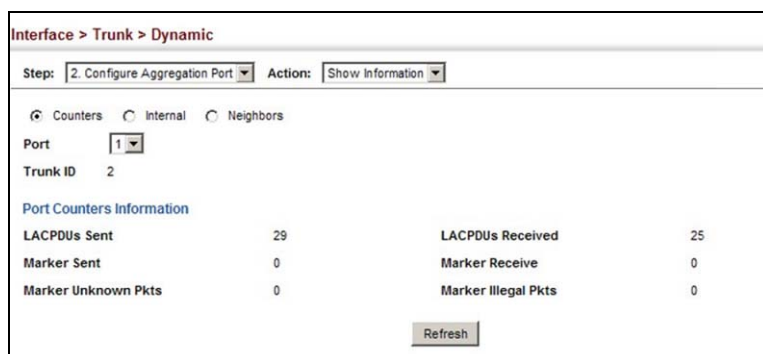


FIG. 74 Displaying LACP Port Counters

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.

The following table lists the options on this page:

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

Perform these steps 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.

The screenshot shows a web-based configuration page titled "Interface > Trunk > Dynamic". At the top, there are two dropdown menus: "Step: 2. Configure Aggregation Port" and "Action: Show Information". Below these are three radio buttons: "Counters", "Internal" (which is selected), and "Neighbors". There is a "Port" dropdown menu set to "1" and a "Trunk ID" field set to "2". Under the heading "Port Internal Information", the following parameters are listed:

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

FIG. 75 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.

The following table lists the options on this page:

LACP Remote Device Configuration Information	
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.)

Perform these steps to display LACP settings and status for the remote side:

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

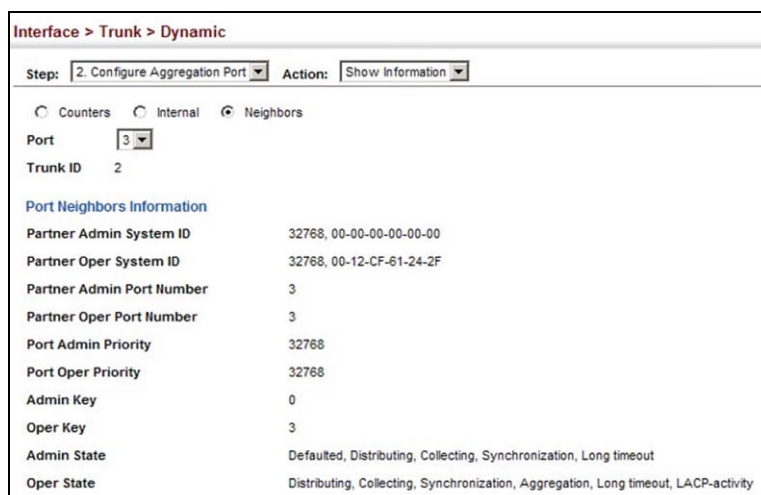


FIG. 76 Displaying LACP Port Remote Information

Configuring Load Balancing

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

Command Usage

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

The following table lists the options on this page:

Interface - Trunk Load Balance Options	
Destination IP Address	Load balancing based on destination IP address.
Destination MAC Address	Load balancing based on destination MAC address.
Source and Destination IP Address	Load balancing based on source and destination IP address.
Source and Destination MAC Address	Load balancing based on source and destination MAC address.
Source IP Address	Load balancing based on source IP address.
Source MAC Address	Load balancing based on source MAC address.

Perform these steps to display the load-distribution method used by ports in aggregated links:

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




FIG. 77 Configuring Load Balancing

Saving Power

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

Command Usage

- IEEE 802.3 defines the Ethernet standard and subsequent power requirements based on cable connections operating at 100 meters. Enabling power saving mode can reduce power used for cable lengths of 60 meters or less, with more significant reduction for cables of 20 meters or less, and continue to ensure signal integrity.

The power-saving methods provided by this switch include:

- Power saving when there is no link partner:

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

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

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

The following table lists the options on this page:

Interface - Trunk Load Balance Options	
Port	Power saving mode only applies to the Gigabit Ethernet ports using copper media. (Range: 1-8/22/24/48)
Power Saving Status	Adjusts the power provided to ports based on the length of the cable used to connect to other devices. Only sufficient power is used to maintain connection requirements. (Default: Enabled on Gigabit Ethernet RJ-45 ports)

Perform these steps to enable power savings:

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

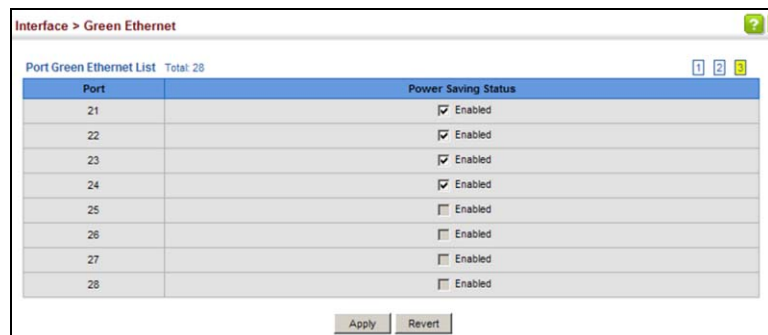


FIG. 78 Enabling Power Savings

Configuring Local Port Mirroring

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

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.
- Monitor port speed should match or exceed source port speed, otherwise traffic may be dropped from the monitor port.
- The destination port cannot be a trunk or trunk member port.
- Note that Spanning Tree BPDU packets are not mirrored to the target port.

The following table lists the options on this page:

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

Perform these steps to configure a local mirror session:

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

FIG. 79 Configuring Local Port Mirroring

Perform these steps to display the configured mirror sessions:

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

	Source (Unit/Port)	Target (Unit/Port)	Type
<input type="checkbox"/>	1 / 7	1 / 8	Both
<input type="checkbox"/>	1 / 9	1 / 10	Both

FIG. 80 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.

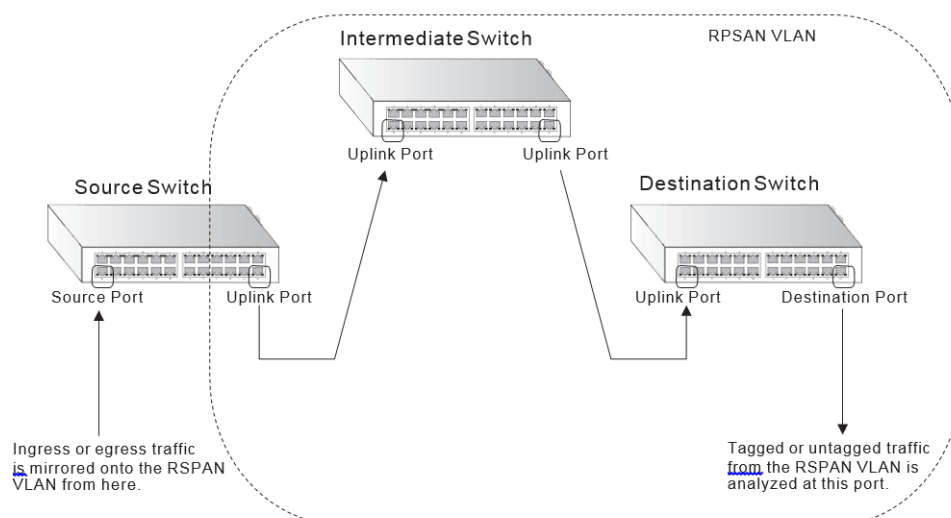


FIG. 81 Configuring Remote Port Mirroring

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 the *Configuring Local Port Mirroring* section on page 81), 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:

- Use the VLAN Static List (see the *Configuring VLAN Groups* section on page 88) to reserve a VLAN for use by RSPAN (marking the "Remote VLAN" field on this page. (Default VLAN 1 is prohibited.)
- Set up the source switch on the RSPAN configuration page by specifying the mirror session, the switch's role (Source), the RSPAN VLAN, and the uplink port*. Then specify the source port(s), and the traffic type to monitor (Rx, Tx or Both).
- 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).
- 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.

* - 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 the *Adding Static Members to VLANs* section on page 89 for more information).

RSPAN Limitations

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

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

The following table lists the options on this page:

Remote Port Mirroring Options	
Session	A number identifying this RSPAN session. (Range: 1-3) Three sessions are allowed, including both local and remote mirroring, using different VLANs for RSPAN sessions.
Operation Status	Indicates whether or not RSPAN is currently functioning.
Switch Role	Specifies the role this switch performs in mirroring traffic. <ul style="list-style-type: none"> • 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 142).
Uplink Port	A port on any switch participating in RSPAN through which mirrored traffic is passed on to or received from the RSPAN VLAN. Only one uplink port can be configured on a source switch, but there is no limitation on the number of uplink ports ¹ 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 port ¹ 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 the traffic exiting the destination port to the monitoring device carries the RSPAN VLAN tag.

Perform these steps 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**.

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

FIG. 82 Configuring Remote Port Mirroring (Source)

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"/>

FIG. 83 Configuring Remote Port Mirroring (Intermediate)

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"/>

FIG. 84 Configuring Remote Port Mirroring (Destination)

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.

The following table lists the options on this page:

Remote Port Mirroring Options	
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. <ul style="list-style-type: none"> Blocking - Blocks traffic between uplink ports assigned to different sessions. Forwarding - Forwards traffic between uplink ports assigned to different sessions.

Perform these steps 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**.

Interface > Traffic Segmentation

Step: 1. Configure Global

Status: Enabled

Uplink-to-Uplink Mode: Blocking

Apply Revert

FIG. 85 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.

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.

- 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.

The following table lists the options on this page:

Traffic Segmentation Options	
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: 1-10/26/28/52)
Trunk	Trunk Identifier (Range: 1-8)

Perform these steps to configure the members of the traffic segmentation group:

- Click **Interface > Traffic Segmentation**.
- Select **Configure Session** from the Step list.
- Select **Add** from the Action list.
- Enter the session ID, set the direction to uplink or downlink, and select the interface to add.
- Click **Apply**.

FIG. 86 Configuring Members for Traffic Segmentation

Perform these steps 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.




FIG. 87 Showing Traffic Segmentation Members

VLAN Configuration

This chapter includes the following topics:

- **IEEE 802.1Q VLANs** - Configures static and dynamic VLANs.
- **Protocol VLANs*** - Configures VLAN groups based on specified protocols.
- **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.

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

IEEE 802.1Q VLANs

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

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

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

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

This switch supports the following VLAN features:

- Up to 4094 VLANs based on the IEEE 802.1Q standard
- Distributed VLAN learning across multiple switches using explicit or implicit tagging 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.*

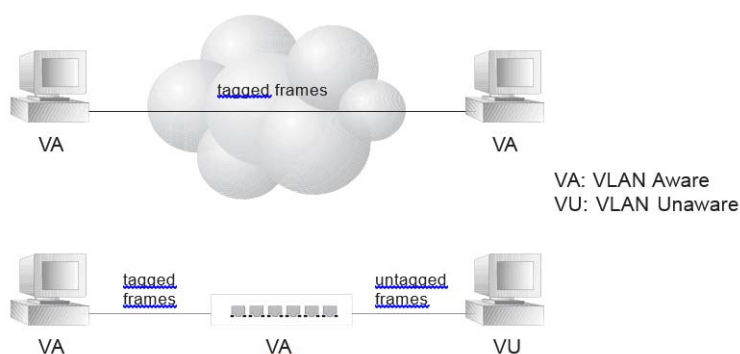


FIG. 88 VLAN Compliant and VLAN Non-compliant Devices

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.

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 the *Configuring Remote Port Mirroring* section on page 82). 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.

The following table lists the options on this page:

VLAN - Static Options	
Add	
VLAN ID	ID of VLAN or range of VLANs (1-4094). VLAN 1 is the default untagged VLAN.
Status	Enables or disables the specified VLAN.
Remote VLAN	Reserves this VLAN for RSPAN (see the <i>Configuring Remote Port Mirroring</i> section on page 82).
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 the <i>Configuring Remote Port Mirroring</i> section on page 82).

Perform these steps to create VLAN groups:

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

The screenshot shows the 'VLAN > Static' configuration page. At the top, the title is 'VLAN > Static'. Below the title, there is a dropdown menu for 'Action:' with 'Add' selected. Underneath, there are three rows of configuration options: 'VLAN ID (1-4094)' with a text input field containing '2', 'Status' with a checked checkbox labeled 'Enabled', and 'Remote VLAN' with an unchecked checkbox labeled 'Enabled'. At the bottom right of the form, there are two buttons: 'Apply' and 'Revert'.

FIG. 89 Creating Static VLANs

Perform these steps to modify the configuration settings for VLAN groups:

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

VLAN > Static

Action: Modify

VLAN ID (1-4094): 1

VLAN Name: DefaultVlan

Status: Enabled

Apply Revert

FIG. 90 Modifying Settings for Static VLANs

Perform these steps to show the configuration settings for VLAN groups:

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

VLAN > Static

Action: Show

Static VLAN List Total: 2

	VLAN ID	VLAN Name	Status	Remote VLAN
<input type="checkbox"/>	1	DefaultVlan	Enabled	Disabled
<input type="checkbox"/>	2	R&D	Enabled	Disabled

Delete Revert

FIG. 91 Showing Static VLANs

Adding Static Members to VLANs

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

The following table lists the options on this page:

VLAN - Static Options	
VLAN	ID of configured VLAN (1-4094).
Interface	Displays a list of ports or trunks.
Port	Port Identifier (Range: 1-10/26/28/52)
Trunk	Trunk Identifier (Range: 1-8)
Mode	Indicates VLAN membership mode for an interface. (Default: Hybrid) <ul style="list-style-type: none"> • Access - Sets the port to operate as an untagged interface. The port transmits and receives untagged frames on a single VLAN only. • Hybrid - Specifies a hybrid VLAN interface. The port may transmit tagged or untagged frames. • 1Q Trunk - Specifies a port as an end-point for a VLAN trunk. A trunk is a direct link between two switches, so the port transmits tagged frames that identify the source VLAN. Note that frames belonging to the port's default VLAN (i.e., associated with the PVID) are also transmitted as tagged frames.
PVID	VLAN ID assigned to untagged frames received on the interface. (Default: 1) When using Access mode, and an interface is assigned to a new VLAN, its PVID is automatically set to the identifier for that VLAN. When using Hybrid mode, the PVID for an interface can be set to any VLAN for which it is an untagged member.
Acceptable Frame Type	Sets the interface to accept all frame types, including tagged or untagged frames, or only tagged frames. When set to receive all frame types, any received frames that are untagged are assigned to the default VLAN. (Options: All, Tagged; Default: All)

VLAN - Static Options	
Ingress Filtering	<p>Determines how to process frames tagged for VLANs for which the ingress port is not a member. (Default: Enabled)</p> <ul style="list-style-type: none"> 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	<p>Select VLAN membership for each interface by marking the appropriate radio button for a port or trunk:</p> <ul style="list-style-type: none"> Tagged: Interface is a member of the VLAN. All packets transmitted by the port will be tagged, that is, carry a tag and therefore carry VLAN or CoS information. Untagged: Interface is a member of the VLAN. All packets transmitted by the port will be untagged, that is, not carry a tag and therefore not carry VLAN or CoS information. Note that an interface must be assigned to at least one group as an untagged port. Forbidden: Interface cannot be included as a member of the VLAN. None: Interface is not a member of the VLAN. Packets associated with this VLAN will not be transmitted by the interface. <p>NOTE: VLAN 1 is the default untagged VLAN containing all ports on the switch using Hybrid mode.</p>
Port Range	Displays a list of ports. (Range: 1-10/28). This option is only available when you select Edit Member by Interface Range from the Action menu.
Trunk Range	Displays a list of ports. (Range: 1-8). This option is only available when you select Edit Member by Interface Range from the Action menu.

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

Perform these steps to configure static members by the VLAN index:

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

The screenshot shows the 'VLAN > Static' configuration page. The 'Action' is set to 'Edit Member by VLAN'. The 'VLAN' is set to '1'. The 'Interface' type is set to 'Port'. Below this is a table titled 'Static VLAN Port Member List' with a total of 28 members. The table has columns for Port, Mode, PVID (1-4094), Acceptable Frame Type, Ingress Filtering, and Membership Type (Tagged, Untagged, Forbidden, None). The first five rows are visible, showing ports 1 through 5, all in Hybrid mode with PVID 1, Acceptable Frame Type All, and Ingress Filtering Enabled. The 'Untagged' membership type is selected for all ports.

Port	Mode	PVID (1-4094)	Acceptable Frame Type	Ingress Filtering	Membership Type			
					Tagged	Untagged	Forbidden	None
1	Hybrid	1	All	Enabled	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
2	Hybrid	1	All	Enabled	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
3	Hybrid	1	All	Enabled	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
4	Hybrid	1	All	Enabled	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
5	Hybrid	1	All	Enabled	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>

FIG. 92 Configuring Static Members by VLAN Index

Perform these steps 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**.

VLAN > Static

Action: Edit Member by Interface

Interface: Port Trunk

Mode: Hybrid

PVID: 1

Acceptable Frame Type: All

Ingress Filtering: Enabled

Static VLAN Membership List Total: 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

FIG. 93 Configuring Static VLAN Members by Interface

Perform these steps 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**.

VLAN > Static

Action: Edit Member by Interface Range

Interface: Port Trunk

Port Range (1-28): [] - []

Mode: Hybrid

VLAN ID (1-4093): [] - []

Membership Type: Tagged Untagged Forbidden None

Apply Revert

FIG. 94 Configuring Static VLAN Members by Interface Range

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 the *Configuring VLAN Groups* section on page 88). 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. Next, map the protocol for each interface to the appropriate VLAN using the Configure Interface (Add) page.

When MAC-based, IP subnet-based, or protocol-based VLANs are supported concurrently, priority is applied in this sequence, and then port-based VLANs last.

Configuring Protocol VLAN Groups

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

The following table lists the options on this page:

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

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

Perform these steps 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**.

The screenshot shows the 'VLAN > Protocol' configuration page. At the top, there are two dropdown menus: 'Step: 1. Configure Protocol' and 'Action: Add'. Below these are three input fields: 'Frame Type' with a dropdown menu set to 'Ethernet', 'Protocol Type' with a dropdown menu set to '08 06 (ARP)', and 'Protocol Group ID (1-2147483647)' with a text input field containing '1'. At the bottom right of the form are two buttons: 'Apply' and 'Revert'.

FIG. 95 Configuring Protocol VLANs

Perform these steps to configure a protocol group:

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

The screenshot shows the 'VLAN > Protocol' configuration page. At the top, there is a 'Step' dropdown set to '1. Configure Protocol' and an 'Action' dropdown set to 'Show'. Below this is a table titled 'Protocol to Group Mapping Table' with a 'Total: 5' indicator. The table has four columns: a checkbox, 'Frame Type', 'Protocol Type', and 'Protocol Group ID'. There are five 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"/>	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

FIG. 96 Displaying Protocol VLANs

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, 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.

The following table lists the options on this page:

VLAN - Protocol Options	
Interface	Displays a list of ports or trunks.
Port	Port Identifier (Range: 1-10/26/28/52)
Trunk	Trunk Identifier (Range: 1-8)
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)

Perform these steps 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**.

The screenshot shows the 'VLAN > Protocol' configuration page at the '2. Configure Interface' step. The 'Action' dropdown is set to 'Add'. There are two radio buttons for 'Interface': 'Port' (selected) and 'Trunk'. Below these are input fields for 'Protocol Group ID' (set to '1'), 'VLAN ID (1-4093)', and 'Priority (0-7)'. At the bottom are 'Apply' and 'Revert' buttons.

FIG. 97 Assigning Interfaces to Protocol VLANs

Perform these steps 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.




FIG. 98 Showing the Interface to Protocol Group Mapping

Configuring MAC-based VLANs

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

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

Command Usage

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

The following table lists the options on this page:

VLAN - MAC-Based Options	
MAC Address	A source MAC address which is to be mapped to a specific VLAN. The MAC address must be specified in the format xx-xx-xx-xx-xx-xx.
Mask	Identifies a range of MAC addresses. (Range: 00-00-00-00-00-00 to ff-ff-ff-ff-ff-ff) The binary equivalent mask matching the characters in the front of the first non-zero character must all be 1s (e.g., 111, i.e., it cannot be 101 or 001...). A mask for the MAC address: 00-50-6e-00-5f-b1 translated into binary: MAC: 00000000-01010000-01101110-00000000-01011111-10110001 could be: 11111111-11xxxxxx-xxxxxxx-xxxxxxx-xxxxxxx-xxxxxxx So the mask in hexadecimal for this example could be: ff-fx-xx-xx-xx-xx/ff-c0-00-00-00-00/ff-e0-00-00-00-00
VLAN	VLAN to which ingress traffic matching the specified source MAC address is forwarded. (Range: 1-4094)
Priority	The priority assigned to untagged ingress traffic. (Range: 0-7, where 7 is the highest priority; Default: 0)

Perform these steps to map a MAC address to a VLAN:

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

VLAN > MAC-Based

Action: Add ▾

MAC Address: 00-ab-cd-11-22-33

Mask:

VLAN (1-4093): 10

Priority (0-7):

Apply Revert

FIG. 99 Configuring MAC-Based VLANs

Perform these steps to show the MAC addresses mapped to a VLAN:

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

VLAN > MAC-Based

Action: Show ▾

MAC-Based VLAN List Total: 1

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

Delete Revert

FIG. 100 Showing MAC-Based VLANs

Address Table Settings

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

This chapter describes the following topics:

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

Configuring MAC Address Learning

Use the MAC Address > Learning Status page to enable or disable MAC address learning on an interface.

Command Usage

- When MAC address learning is disabled, the switch immediately stops learning new MAC addresses on the specified interface. Only incoming traffic with source addresses stored in the static address table (see the *Setting Static Addresses* section on page 97) will be accepted as authorized to access the network through that interface.
- Dynamic addresses stored in the address table when MAC address learning is disabled are flushed from the system, and no dynamic addresses are subsequently learned until MAC address learning has been re-enabled. Any device not listed in the static address table that attempts to use the interface after MAC learning has been disabled will be prevented from accessing the switch.
- Also note that MAC address learning cannot be disabled if any of the following conditions exist:
 - 802.1X Port Authentication has been globally enabled on the switch (see the *Configuring 802.1x Port Authentication* section on page 173).
 - Security Status (see the *Configuring Port Security* section on page 172) is enabled on the same interface.

The following table lists the options on this page:

MAC Address - Learning Status Options	
Interface	Displays a list of ports or trunks.
Port	Port Identifier (Range: 1-10/28)
Trunk	Trunk Identifier (Range: 1-8)
Status	The status of MAC address learning. (Default: Enabled)

Perform these steps to enable or disable MAC address learning:

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

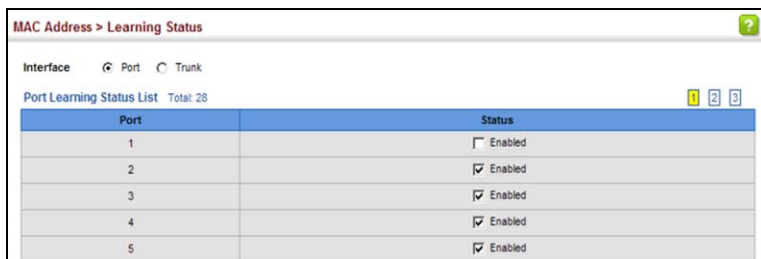


FIG. 101 Configuring MAC Address Learning

Setting Static Addresses

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

Command Usage

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

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

The following table lists the options on this page:

MAC Address - Static Options	
Add Static Address	
VLAN	ID of configured VLAN (Range: 1-4094)
Interface	Port or trunk associated with the device assigned a static address.
MAC Address	Physical address of a device mapped to this interface. Enter an address in the form of xx-xx-xx-xx-xx-xx or xxxxxxxxxxxx.
Static Status	Sets the time to retain the specified address. <ul style="list-style-type: none"> • Delete-on-reset - Assignment lasts until the switch is reset. • Permanent - Assignment is permanent. (This is the default.)
Show Static Address	
Type	Displays the address configuration method. (Values: CPU, Config, or Security, the last of which indicates Port Security)
Life Time	The duration for which this entry applies. (Values: Delete On Reset, Delete On Timeout, Permanent)

Perform these steps to configure a static MAC address:

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

FIG. 102 Configuring Static MAC Addresses

Perform these steps to show the static addresses in MAC address table:

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

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

FIG. 103 Displaying Static MAC Addresses

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.

The following table lists the options on this page:

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

Perform these steps 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**.

The screenshot shows a web interface for configuring aging settings. At the top, it says 'MAC Address > Dynamic'. Below that, there's a dropdown menu for 'Action' with 'Configure Aging' selected. Underneath, there's a section for 'Aging Status' with a checked checkbox and the text 'Enabled'. Below that, there's a section for 'Aging Time (6-7200)' with a text input field containing '300' and the unit 'sec'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

FIG. 104 Setting the Address Aging Time

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.

The following table lists the options on this page:

MAC Address - Dynamic Options	
Sort Key	You can sort the information displayed based on MAC address, VLAN or interface (port or trunk).
MAC Address	Physical address associated with this interface.
VLAN	ID of configured VLAN (1-4094).
Interface	Indicates a port or trunk.
Type	Shows that the entries in this table are learned. (Values: Learned or Security, the last of which indicates Port Security)
Life Time	Shows the time to retain the specified address.

Perform these steps 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**.

The screenshot shows the 'MAC Address > Dynamic' configuration page. The 'Action' dropdown is set to 'Show Dynamic MAC'. Under 'Query by:', the 'Sort Key' is 'MAC Address'. There are checkboxes for 'MAC Address', 'VLAN', and 'Interface'. The 'Interface' section has radio buttons for 'Port' and 'Trunk', with 'Port' selected and '1' entered in the adjacent field. A 'Query' button is visible. Below is a table titled 'Dynamic MAC Address List' with a total of 2 entries.

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

FIG. 105 Displaying the Dynamic MAC Address Table

Clearing the Dynamic Address Table

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

The following table lists the options on this page:

MAC Address - Dynamic Options	
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.

Perform these steps 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**.

The screenshot shows the 'MAC Address > Dynamic' configuration page. The 'Step' dropdown is set to '3. Clear Dynamic MAC'. The 'Clear by:' dropdown is set to 'All'. A 'Clear' button is visible at the bottom right.

FIG. 106 Clearing Entries in the Dynamic MAC Address Table

Issuing MAC Address Traps

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

The following table lists the options on this page:

MAC Address - MAC Notification Options	
Configure Global	
MAC Notification Traps	Issues a trap when a dynamic MAC address is added or removed. (Default: Disabled)
MAC Notification Trap Interval	Specifies the interval between issuing two consecutive traps. (Range: 1-3600 seconds; Default: 1 second)
Configure Interface	
Port	Port Identifier (Range: 1-10/26/28/52)
MAC Notification Trap	Enables MAC authentication traps on the current interface. (Default: Disabled) MAC authentication traps must be enabled at the global level for this attribute to take effect.

Perform these steps to enable MAC address traps at the global level:

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

FIG. 107 Issuing MAC Address Traps (Global Configuration)

Perform these steps to enable MAC address traps at the interface level:

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

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

FIG. 108 Issuing MAC Address Traps (Interface Configuration)

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.




FIG. 109 STP Root Ports and Designated Ports

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

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

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




FIG. 110 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 the *Configuring Multiple Spanning Trees* section on page 113). 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.

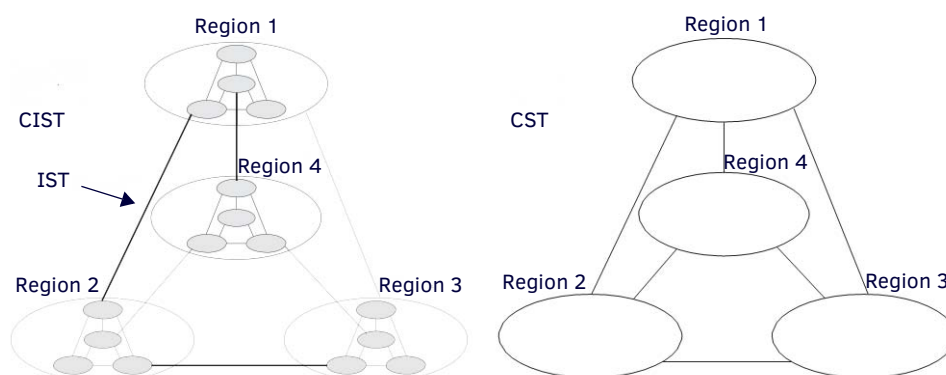


FIG. 111 Spanning Tree - Common Internal, Common, Internal

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

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

Configuring Loopback Detection

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

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

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

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

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

The following table lists the options on this page:

Spanning Tree - Loopback Detection Options	
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 shut down the interface. (Default: Shutdown)
Shutdown Interval	The duration to shut down the interface. (Range: 60-86400 seconds; Default: 60 seconds) If an interface is shut down due to a detected loopback, and the release mode is set to Auto, the selected interface will be automatically enabled when the shutdown interval has expired. If an interface is shut down due to a detected loopback, and the release mode is set to Manual, the interface can be re-enabled using the Release button.

Perform these steps 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**

Port	Status	Trap	Release Mode	Release	Action	Shutdown Interval (60-86400 sec)
1	<input type="checkbox"/> Enabled	<input checked="" type="checkbox"/> Enabled	Auto	Release	Shutdown	60
2	<input type="checkbox"/> Enabled	<input checked="" type="checkbox"/> Enabled	Auto	Release	Shutdown	60
3	<input type="checkbox"/> Enabled	<input checked="" type="checkbox"/> Enabled	Auto	Release	Shutdown	60
4	<input type="checkbox"/> Enabled	<input checked="" type="checkbox"/> Enabled	Manual	Release	Shutdown	60
5	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled	Auto	Release	Shutdown	60

FIG. 112 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 Protocol*** - 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 Protocol*** - 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.

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

The following table lists the options on this page:

Spanning Tree - STA Options	
Basic Configuration of Global Settings	
Spanning Tree Status	Enables/disables STA on this switch. (Default: Disabled) When spanning tree is enabled globally or enabled on an interface, loopback detection is disabled.
Spanning Tree Type	Specifies the type of spanning tree used on this switch: <ul style="list-style-type: none"> • STP: Spanning Tree Protocol (IEEE 802.1D); i.e., when this option is selected, the switch will use RSTP set to STP forced compatibility mode). • RSTP: Rapid Spanning Tree (IEEE 802.1w); RSTP is the default. • MSTP: Multiple Spanning Tree (IEEE 802.1s)
Priority	Bridge priority is used in selecting the root device, root port, and designated port. The device with the highest priority becomes the STA root device. However, if all devices have the same priority, the device with the lowest MAC address will then become the root device. (Note that lower numeric values indicate higher priority.) <ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • 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.
Cisco Prestandard Status	Configures spanning tree operation to be compatible with Cisco pre-standard versions. (Default: Disabled) Cisco pre-standard versions prior to Cisco IOS Release 12.2(25)SEC do not fully follow the IEEE standard, causing some state machine procedures to function incorrectly. This command forces the spanning tree protocol to function in a manner compatible with Cisco prestandard versions.
Advanced Configuration Settings	
The following attributes are based on RSTP, but also apply to STP since the switch uses a backwards-compatible subset of RSTP to implement STP, and also apply to MSTP which is based on RSTP according to the standard.	
Path Cost Method	The path cost is used to determine the best path between devices. The path cost method is used to determine the range of values that can be assigned to each interface. <ul style="list-style-type: none"> • 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)

Spanning Tree - STA Options	
When the Switch Becomes Root	
Hello Time	Interval (in seconds) at which the root device transmits a configuration message. <ul style="list-style-type: none"> • 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.) <ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • Default: 15 • Minimum: The higher of 4 or $[(\text{Max. Message Age} / 2) + 1]$ • Maximum: 30 <p>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.</p>
Configuration Settings for MSTP	
Max Instance Numbers	The maximum number of MSTP instances to which this switch can be assigned.
Configuration Digest	An MD5 signature key that contains the VLAN ID to MST ID mapping table. In other words, this key is a mapping of all VLANs to the CIST.
Region Revision	The revision for this MSTI. (Range: 0-65535; Default: 0). Region Revision is required to uniquely identify an MST region.
Region Name	The name for this MSTI. (Maximum length: 32 characters; Default: switch's MAC address). Region Name is required to uniquely identify an MST region.
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.*

Perform these steps 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**.

The screenshot shows the configuration page for Spanning Tree > STA. The 'Step' is set to '1. Configure Global' and the 'Action' is 'Configure'. The 'Spanning Tree Status' is checked and 'Enabled'. The 'Spanning Tree Type' is set to 'STP'. The 'Priority (0-61440, in steps of 4096)' is 32768. 'BPDU Flooding' is set to 'To VLAN'. 'Cisco Prestandard Status' is unchecked. Under 'Advanced', 'Path Cost Method' is 'Long' and 'Transmission Limit (1-10)' is 3. Under 'When the Switch Becomes Root', 'Hello Time (1-10)' is 2 sec, 'Maximum Age (6-40)' is 20 sec, and 'Forward Delay (4-30)' is 15 sec. A note at the bottom states: $2 * (\text{Hello Time} + 1) \leq \text{Max Age} \leq 2 * (\text{Forward Delay} - 1)$. 'Apply' and 'Revert' buttons are at the bottom right.

FIG. 113 Configuring Global Settings for STA (STP)

The screenshot shows the configuration page for Spanning Tree > STA. The 'Step' is set to '1. Configure Global' and the 'Action' is 'Configure'. The 'Spanning Tree Status' is checked and 'Enabled'. The 'Spanning Tree Type' is set to 'RSTP'. The 'Priority (0-61440, in steps of 4096)' is 32768. 'BPDU Flooding' is set to 'To VLAN'. 'Cisco Prestandard Status' is unchecked. Under 'Advanced', 'Path Cost Method' is 'Long' and 'Transmission Limit (1-10)' is 3. Under 'When the Switch Becomes Root', 'Hello Time (1-10)' is 2 sec, 'Maximum Age (6-40)' is 20 sec, and 'Forward Delay (4-30)' is 15 sec. A note at the bottom states: $2 * (\text{Hello Time} + 1) \leq \text{Max Age} \leq 2 * (\text{Forward Delay} - 1)$. 'Apply' and 'Revert' buttons are at the bottom right.

FIG. 114 Configuring Global Settings for STA (RSTP)

Spanning Tree > STA

Step: 1. Configure Global Action: Configure

Spanning Tree Status Enabled

Spanning Tree Type MSTP

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

BPDU Flooding To VLAN

Cisco Prestandard Status Enabled

Advanced:

Path Cost Method Long

Transmission Limit (1-10) 3

When the Switch Becomes Root:

Hello Time (1-10) 2 sec

Maximum Age (6-40) 20 sec

Forward Delay (4-30) 15 sec

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

MSTP Configuration

Max Instance Numbers 64

Configuration Digest 0xAC36177F50283CD4B83821D8AB26DE62

Region Revision (0-65535) 0

Region Name 00 E0 0C 00 00 FD

Max Hop Count (1-40) 20

Apply Revert

FIG. 115 Configuring Global Settings for STA (MSTP)

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.

The following table lists the options on this page:

Spanning Tree - STA Options	
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.

Perform these steps 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.

Spanning Tree > STA

Step: 1. Configure Global Action: Show Information

Spanning Tree Information

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

FIG. 116 Displaying Global Settings for STA

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.)

The following table lists the options on this page:

Spanning Tree - STA Options	
Interface	Displays a list of ports or trunks.
Spanning Tree	Enables/disables STA on this interface. (Default: Enabled) When spanning tree is enabled globally (Configuring Global Settings for STA) or enabled on an interface by this command, loopback detection is disabled.
BPDU Flooding	Enables/disables the flooding of BPDUs to other ports when global spanning tree is disabled (page 169) 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 169). (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. <ul style="list-style-type: none"> • 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 method*, 1-200,000,000 for the long path cost method) By default, the system automatically detects the speed and duplex mode used on each port, and configures the path cost according to the values shown below. Path cost "0" is used to indicate auto-configuration mode. When the short path cost method is selected and the default path cost recommended by the IEEE 8021w standard exceeds 65,535, the default is set to 65,535. See the Recommended STA Path Cost Range table on page 109 and the Default STA Path Costs table on page 110 for more information. * - Refer to the <i>Configuring Global Settings for STA</i> section on page 104 for information on setting the path cost method. The range displayed on the STA interface configuration page shows the maximum value for path cost. However, note that the switch still enforces the rules for path cost based on the specified path cost method (long or short)
Admin Link Type	The link type attached to this interface. <ul style="list-style-type: none"> • 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)

Spanning Tree - STA Options	
Admin Edge Port	<p>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)</p> <ul style="list-style-type: none"> • 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 the <i>Configuring Global Settings for STA</i> section on page 104). <p>An interface cannot function as an edge port under the following conditions:</p> <ul style="list-style-type: none"> • If spanning tree mode is set to STP, edge-port mode cannot automatically transition to operational edge-port state using the automatic setting. • If loopback detection is enabled 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 the <i>Displaying Interface Settings for STA</i> section on page 111). <p>When edge port is set as auto, the operational state is determined automatically by the Bridge Detection State Machine described in 802.1D-2004, where the edge port state may change dynamically based on environment changes (e.g., receiving a BPDU or not within the required interval).</p>
BPDU Guard	<p>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)</p> <p>BPDU guard can only be configured on an interface if the edge port attribute is not disabled (that is, if edge port is set to enabled or auto).</p>
BPDU Guard Auto Recovery	Automatically re-enables an interface after the specified interval. (Range: 30-86400 seconds; Default: Disabled)
BPDU Guard Auto Recovery Interval	The time to wait before re-enabling an interface. (Range: 30-86400 seconds; Default: 300 seconds)
BPDU Filter	<p>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)</p> <p>BPDU filter can only be configured on an interface if the edge port attribute is not disabled (that is, if edge port is set to enabled or auto).</p>
Migration	If at any time the switch detects STP BPDUs, including Configuration or Topology Change Notification BPDUs, it will automatically set the selected interface to forced STP-compatible mode. However, you can also use the Protocol Migration button to manually re-check the appropriate BPDU format (RSTP or STP-compatible) to send on the selected interfaces. (Default: Disabled)
TC Propagate Stop	Stops the propagation of topology change notifications (TCN). (Default: Disabled)

Recommended STA Path Cost Range		
Port Type	IEEE 802.1D-1998	IEEE 802.1w-2001
Ethernet	50-600	200,000-20,000,000
Fast Ethernet	10-60	20,000-2,000,000
Gigabit Ethernet	3-10	2,000-200,000
10G Ethernet	1-5	200-20,000

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

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




FIG. 117 Determining the Root Port

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

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

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

If the path cost of i1 on SW2 is never configured/changed, it is 10000. Then the root path cost for i2 on SW3 used to compete for the role of root port is 10000 + path cost of i2 on SW3.

The path cost of i1 on SW3 is also 10000 if not configured/changed. Then even if the path cost of i2 on SW3 is configured/changed to 0, these ports will still have the same root path cost, and it will be impossible for i2 to become the root port just by changing its path cost on SW3.

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

Perform these steps 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**.

Port	Spanning Tree	BPDU Flooding	Priority (0-240, in steps of 16)	Admin Path Cost (0-200000000, 0: Auto)	Admin Link Type	Root Guard	Admin Edge Port	BPDU Guard	BPDU Guard Auto Recovery	BPDU Guard Auto Recovery Interval (30-96400)	BPDU Filter	Migration	TC Propagate Stop
1	Enabled	Enabled	128	0	Auto	Enabled	Auto	Enabled	Enabled	300	Enabled	Enabled	Enabled
2	Enabled	Enabled	128	0	Auto	Enabled	Auto	Enabled	Enabled	300	Enabled	Enabled	Enabled
3	Enabled	Enabled	128	0	Auto	Enabled	Auto	Enabled	Enabled	300	Enabled	Enabled	Enabled
4	Enabled	Enabled	128	0	Auto	Enabled	Auto	Enabled	Enabled	300	Enabled	Enabled	Enabled
5	Enabled	Enabled	128	0	Auto	Enabled	Auto	Enabled	Enabled	300	Enabled	Enabled	Enabled

FIG. 118 Configuring Interface Settings for STA

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.

The following table lists the options on this page:

Spanning Tree - STA Options	
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	<p>Displays current state of this port within the Spanning Tree:</p> <ul style="list-style-type: none"> 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. <p>The rules defining port status are:</p> <ul style="list-style-type: none"> 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 <i>Admin Link Type</i> in the Spanning Tree - STA Options table on page 108.
Oper Edge Port	This parameter is initialized to the setting for <i>Admin Edge Port</i> in the Spanning Tree - STA Options table on page 108 (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.

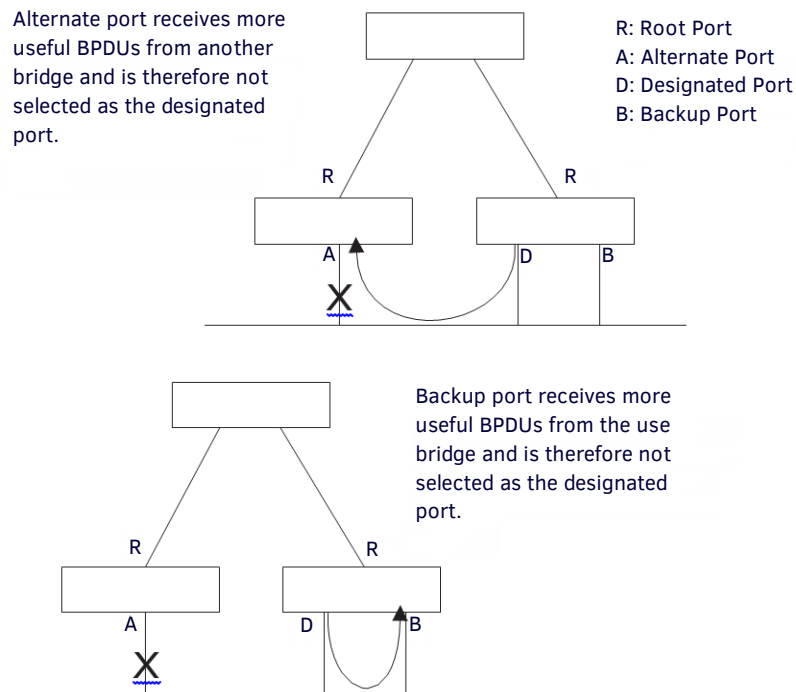


FIG. 119 STA Port Roles

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

Perform these steps 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.

Spanning Tree > STA

Step: 2. Configure Interface Action: Show Information

Interface Port Trunk

Spanning Tree Port List Total: 28

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

FIG. 120 Displaying Interface Settings for STA

Configuring Multiple Spanning Trees

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

Command Usage

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

By default all VLANs are assigned to the Internal Spanning Tree (MST Instance 0) that connects all bridges and LANs within the MST region. This switch supports up to 33 instances. You should try to group VLANs which cover the same general area of your network. However, remember that you must configure all bridges within the same MSTI Region (page 104) 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.

Perform these steps to use multiple spanning trees:

1. Set the spanning tree type to MSTP (page 104).
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.

The following table lists the options on this page:

Spanning Tree - MSTP Options	
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)

Perform these steps 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**.

FIG. 121 Creating an MST Instance

Perform these steps 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.

FIG. 122 Displaying MST Instances

Perform these steps 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**.

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

FIG. 123 Modifying the Priority for an MST Instance

Perform these steps 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 the *Displaying Global Settings for STA* section on page 107.

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

FIG. 124 Displaying Global Settings for an MST Instance

Perform these steps 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**.

FIG. 125 Adding a VLAN to an MST Instance

Perform these steps 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.

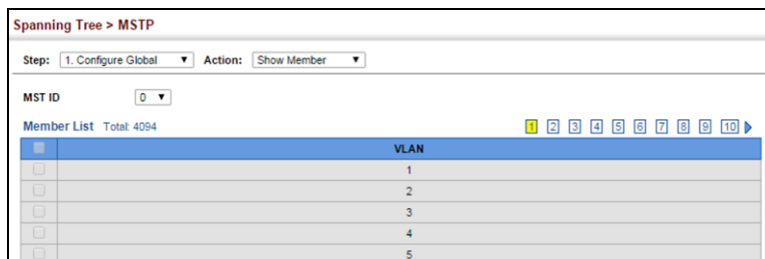


FIG. 126 Displaying Members of an MST Instance

Configuring Interface Settings for MSTP

Use the Spanning Tree > MSTP (Configure Interface - Configure) page to configure the STA interface settings for an MST instance. The following table lists the options on this page:

Spanning Tree - MSTP Options	
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 the "Displaying Interface Settings for STA" section on page 111 for more information.) <ul style="list-style-type: none"> • 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 104), 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 802.1w standard exceeds 65,535, the default is set to 65,535. The recommended range is listed in the Recommended STA Path Cost Range table on page 109. The default path costs are listed in the Default STA Path Costs table on page 110.

Perform these steps 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**.

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

FIG. 127 Configuring MSTP Interface Settings

Perform these steps 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.

Spanning Tree > MSTP

Step: 2. Configure Interface Action: Show Information

MST ID: 0

Interface: Port Trunk

Spanning Tree Port List Total: 28

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

FIG. 128 Displaying MSTP Interface Settings

Congestion Control

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

Congestion Control includes following options:

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

Rate Limiting

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

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

The following table lists the options on this page:

Traffic - Rate Limit Options	
Interface	Displays the switch's ports or trunks.
Type	Indicates the port type (1000BASE-T, 10GBASE SFP+, or 1000BASE SFP (used when the transceiver type is used in an SFP+ port))
Status	Enables or disables the rate limit. (Default: Disabled)
Rate	Sets the rate limit level. (Range: 64 - 1,000,000 kbits per second for Gigabit Ethernet ports; 64 - 10,000,000 kbits per second for 10 Gigabit Ethernet ports)
Resolution	Indicates the resolution at which the rate can be configured.

Perform these steps to configure rate limits:

1. Click **Traffic > Rate Limit**.
2. Set the interface type to Port or Trunk.
3. Enable the **Rate Limit Status** for the required interface.
4. Set the rate limit for required interfaces.
5. Click **Apply**.

The screenshot shows the 'Traffic > Rate Limit' configuration page. At the top, there are radio buttons for 'Interface' with 'Port' selected and 'Trunk' unselected. Below this is a 'Port Rate Limit List' with a 'Total: 28' and three icons (list, refresh, delete). The main table has columns for Port, Type, Status, Input Rate (kbits/sec), Output Rate (kbits/sec), and Resolution (kbits/sec). The table contains 5 rows of data, all with '1000BASE-T' type and 'Enabled' status, with input and output rates set to 1000000 and a resolution of 16.

Port	Type	Status	Input		Output		Resolution (kbits/sec)
			Status	Rate (kbits/sec) (64-1000000)	Status	Rate (kbits/sec) (64-1000000)	
1	1000BASE-T	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> Enabled	1000000	<input checked="" type="checkbox"/> Enabled	1000000	16
2	1000BASE-T	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> Enabled	1000000	<input checked="" type="checkbox"/> Enabled	1000000	16
3	1000BASE-T	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> Enabled	1000000	<input checked="" type="checkbox"/> Enabled	1000000	16
4	1000BASE-T	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> Enabled	1000000	<input checked="" type="checkbox"/> Enabled	1000000	16
5	1000BASE-T	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/> Enabled	1000000	<input checked="" type="checkbox"/> Enabled	1000000	16

FIG. 129 Configuring Rate Limits

Storm Control

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

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

Command Usage

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

The following table lists the options on this page:

Traffic - Rate Limit Options	
Interface	Displays a list of ports or trunks.
Type	Indicates the port type. (1000BASE-T, 10GBASE SFP+, or 1000BASE SFP (used in the ECS4620-28F/28F-DC or when this transceiver type is used in an SFP+ port).
Unknown Unicast	Specifies storm control for unknown unicast traffic.
Multicast	Specifies storm control for multicast traffic.
Broadcast	Indicates the resolution at which the rate can be configured.
Status	Enables or disables storm control. (Default: Disabled)
Rate	Threshold level in packets per second. (Range: 500-262142 pps; Default: 500 pps)
Resolution	Indicates the resolution at which the rate can be configured.

Perform these steps 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**.

Port	Type	Unknown Unicast		Multicast		Broadcast		Resolution (packets/sec)
		Status	Rate (packets/sec) (500-262142)	Status	Rate (packets/sec) (500-262142)	Status	Rate (packets/sec) (500-262142)	
1	1000BASE-T	<input checked="" type="checkbox"/> Enabled	500	<input checked="" type="checkbox"/> Enabled	500	<input checked="" type="checkbox"/> Enabled	500	1
2	1000BASE-T	<input checked="" type="checkbox"/> Enabled	500	<input checked="" type="checkbox"/> Enabled	500	<input checked="" type="checkbox"/> Enabled	500	1
3	1000BASE-T	<input checked="" type="checkbox"/> Enabled	500	<input checked="" type="checkbox"/> Enabled	500	<input checked="" type="checkbox"/> Enabled	500	1
4	1000BASE-T	<input checked="" type="checkbox"/> Enabled	500	<input checked="" type="checkbox"/> Enabled	500	<input checked="" type="checkbox"/> Enabled	500	1
5	1000BASE-T	<input checked="" type="checkbox"/> Enabled	500	<input checked="" type="checkbox"/> Enabled	500	<input checked="" type="checkbox"/> Enabled	500	1

FIG. 130 Configuring Storm Control

Class of Service

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

This chapter describes the following basic topics:

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

Layer 2 Queue Settings

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

Setting the Default Priority for Interfaces

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

Command Usage

- This switch provides eight priority queues for each port. It uses Weighted Round Robin to prevent head-of-queue blockage, but can be configured to process each queue in strict order, or use a combination of strict and weighted queuing.
- 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.

The following table lists the options on this page:

Traffic - Priority (Default Priority) Options	
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)

Perform these steps 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**.

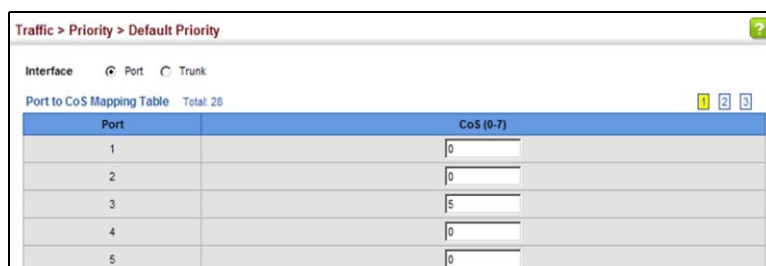


FIG. 131 Setting the Default Port Priority

Selecting the Queue Mode

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

Command Usage

- Strict priority requires all traffic in a higher priority queue to be processed before lower priority queues are serviced.
- WRR queuing specifies a relative weight for each queue. WRR uses a predefined relative weight for each queue that determines the percentage of service time the switch services each queue before moving on to the next queue. This prevents the head-of-line blocking that can occur with strict priority queuing.
- If Strict and WRR mode is selected, a combination of strict service is used for the high priority queues and weighted service for the remaining queues. The queues assigned to use strict priority should be specified using the Strict Mode field parameter.
- A weight can be assigned to each of the weighted queues (and thereby to the corresponding traffic priorities). This weight sets the frequency at which each queue is polled for service, and subsequently affects the response time for software applications assigned a specific priority value.

Service time is shared at the egress ports by defining scheduling weights for WRR, or one of the queuing modes that use a combination of strict and weighted queuing.

- The specified queue mode applies to all interfaces.

The following table lists the options on this page:

Traffic - Priority (Queue) Options	
Queue Mode	<ul style="list-style-type: none"> • Strict - Services the egress queues in sequential order, transmitting all traffic in the higher priority queues before servicing lower priority queues. This ensures that the highest priority packets are always serviced first, ahead of all other traffic. • WRR - Weighted Round-Robin shares bandwidth at the egress ports by using scheduling weights, and servicing each queue in a round-robin fashion. (This is the default setting.) • Strict and WRR - Uses strict priority on the high-priority queues and WRR on the remaining queues.
Queue ID	The ID of the priority queue. (Range: 0-7)
Strict Mode	If Strict and WRR mode is selected, then a combination of strict service is used for the high priority queues and weighted service for the remaining queues. Use this parameter to specify the queues assigned to use strict priority when using the strict-weighted queuing mode. (Default: Disabled)
Weight	Sets a weight for each queue which is used by the WRR scheduler. (Range: 1-127; Default: Weights 1, 2, 4, 6, 8, 10, 12 and 14 are assigned to queues 0 - 7 respectively)

Perform these steps 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**.

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

FIG. 132 Setting the Queue Mode (Strict)

Traffic > Priority > Queue

Port: 1

Queue Mode: WRR

Queue Setting Table Total: 8

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

Apply Revert

FIG. 133 Setting the Queue Mode (WRR)

Traffic > Priority > Queue

Port: 1

Queue Mode: Strict and WRR

Queue Setting Table Total: 8

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

Apply Revert

FIG. 134 Setting the Queue Mode (Strict and WRR)

Layer 3/4 Priority Settings

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 the *Setting the Default Priority for Interfaces* section on page 119) 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 the *Setting the Default Priority for Interfaces* section on page 119) is used for priority processing.

The following table lists the options on this page:

Traffic - Priority (Trust Mode) Options	
Port	Port identifier. (Range: 1-10/28)
Trust Mode	<ul style="list-style-type: none"> • 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.

Perform these steps to configure the trust mode:

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



FIG. 135 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 121), 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.

The following table lists the options on this page:

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

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.

Perform these steps 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 the port to configure.
4. Set the PHB and drop precedence for any DSCP value.
5. Click **Apply**.

FIG. 136 Configuring DSCP to DSCP Internal Mapping

Perform these steps to show the DSCP to internal PHB/drop precedence map:

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

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

FIG. 137 Showing DSCP to DSCP Internal Mapping

Mapping CoS Priorities to Internal DSCP Values

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

Command Usage

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

The following table lists the options on this page:

Traffic - Priority (CoS to DSCP) Options	
Port	Specifies a port. (Range: 1-10/26/28/52)
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)

Default Mapping of CoS/CFI to Internal PHB/Drop Precedence			
		CFI	
		0	1
CoS	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)

Perform these steps 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. Set the PHB and drop precedence for any of the CoS/CFI combinations.
4. Click **Apply**.

FIG. 138 Configuring CoS to DSCP Internal Mapping

Perform these steps 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.

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

FIG. 139 Showing CoS to DSCP Internal Mapping

Quality of Service

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

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

Overview

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

All switches or routers that access the Internet rely on class information to provide the same forwarding treatment to packets in the same class. Class information can be assigned by end hosts, or switches or routers along the path. Priority can then be assigned based on a general policy, or a detailed examination of the packet.

However, note that detailed examination of packets should take place close to the network edge so that core switches and routers are not overloaded.

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

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

NOTE: You should create a class map before creating a policy map. Otherwise, you will not be able to select a class map from the policy rule settings screen (see page 210).

Command Usage

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

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

NOTE: Up to 16 classes can be included in a policy map.

Configuring a Class Map

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

Command Usage

- The class map is used with a policy map (page 128) to create a service policy (page 130) 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.

The following table lists the options on this page:

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

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

Perform these steps to configure a class map:

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

Traffic > DiffServ

Step: 1. Configure Class Action: Add

Class Name: rd-class

Type: Match Any

Description: class for software group

Apply Revert

FIG. 140 Configuring a Class Map

Perform these steps 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.

Traffic > DiffServ

Step: 1. Configure Class Action: Show

Class List Total: 1

	Class Name	Type	Description
<input type="checkbox"/>	rd-class	Match Any	class for software group

Delete Revert

FIG. 141 Showing Class Maps

Perform these steps to edit the rules for a class map:

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

The screenshot shows the 'Traffic > DiffServ' configuration page. The 'Step' is '2. Configure Policy' and the 'Action' is 'Add Rule'. The 'Policy Name' is 'rd-policy'. Under the 'Rule' section, the 'Class Name' is 'rd-class'. The 'Action' is set to 'Set' with a dropdown for 'CoS (0-7)'. The 'Meter' section is expanded, showing 'Meter Mode' as 'Flow'. Below this, several fields are visible for 'Committed Information Rate', 'Committed Burst Size', 'Excess Burst Size', 'Peak Information Rate', and 'Peak Burst Size', each with a text input and a unit dropdown (kbps or bytes). The 'Conform' action is 'Transmit'. The 'Exceed' and 'Violate' actions are both set to 'Set IP DSCP (0-63)' with associated text input fields.

FIG. 142 Adding Rules to a Class Map

Perform these steps 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.

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

Rule List	Total: 2
<input type="checkbox"/>	Rule
<input type="checkbox"/>	IP DSCP 3
<input type="checkbox"/>	IP Precedence 3

FIG. 143 Showing the Rules for a Class Map

Creating QoS Policies

Use the Traffic > DiffServ (Configure Policy) page to create a policy map that can be attached to multiple interfaces. A policy map is used to group one or more class map statements (page 125). A policy map can then be bound by a service policy to one or more interfaces (page 130).

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 a specific VLAN. A policy map is then configured which indicates the boundary parameters used for monitoring inbound 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.

Meter Mode - Defines the committed information rate (maximum throughput).

- Policing is based on a token bucket, where bucket depth is the maximum burst before the bucket overflows, and the average rate tokens that are added to the bucket is by specified by the committed-rate option. Note that the token bucket functions similar to that described in RFC 2697 and RFC 2698.
- 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 $Tc(0) = BC$ and the token count $Te(0) = BE$. Thereafter, the token counts Tc and Te are updated CIR times per second as follows:
 - If Tc is less than BC, Tc is incremented by one, else
 - if Te is less than BE, Te is incremented by one, else
 - neither Tc nor Te 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 $Tc(t) - B \leq 0$, the packet is green and Tc is decremented by B down to the minimum value of 0, else
- if $Te(t) - B \leq 0$, the packets is yellow and Te is decremented by B down to the minimum value of 0,
- else the packet is red and neither Tc nor Te 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 $Tc(t) - B \leq 0$, the packet is green and Tc is decremented by B down to the minimum value of 0, else
- If the packet has been precolored as yellow or green and if $Te(t) - B \leq 0$, the packets is yellow and Te is decremented by B down to the minimum value of 0, else
- the packet is red and neither Tc nor Te is decremented.

The following table lists the options on this page:

Traffic - Diffserv Options	
Add	
Policy Name	Name of policy map. (Range: 1-32 characters)
Description	A brief description of a policy map. (Range: 1-64 characters)
Add Rule	
Policy Name	Name of policy map.
Class Name	Name of a class map that defines a traffic classification upon which a policy can act. A policy map can contain up to 32 class maps.
Action	This attribute is used to set an internal QoS value in hardware for matching packets. The PHB label is composed of five bits, three bits for per-hop behavior, and two bits for the color scheme used to control queue congestion with the srTCM and trTCM metering functions. <ul style="list-style-type: none"> • 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 the Default Mapping of CoS/CFI to Internal PHB/Drop Precedence table on page 124. • 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 the Default Mapping of DSCP Values to Internal PHB/Drop Values table on page 122. • Set IP DSCP - Configures the service provided to ingress traffic by setting an IP DSCP value for a matching packet (as specified in rule settings for a class map). (Range: 0-63)
Meter	Check this to define the maximum throughput. <ul style="list-style-type: none"> • Meter Mode (Rate) - Defines the committed information rate. Policing is based on a token bucket, where the average rate tokens that are removed from the bucket is specified by the Rate option. The committed rate is in kilobits per second. (Range: 16-1000000 kbps at a granularity of 64 kbps or maximum port speed, whichever is lower)

Perform these steps 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 **Apply**.

Traffic > DiffServ

Step: 2. Configure Policy Action: Add

Policy Name: rd-policy

Description: for the software group

Apply Revert

FIG. 144 Configuring a Policy Map

Perform these steps 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.

Traffic > DiffServ

Step: 2. Configure Policy Action: Show

Policy List Total: 1

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

Delete Revert

FIG. 145 Showing Policy Maps

Perform these steps 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. Click on the Action field, and set the CoS or per-hop behavior for matching packets to specify the quality of service to be assigned to the matching traffic class.
6. Use the metering option to define the maximum throughput.
7. Click **Apply**.

Traffic > DiffServ

Step: 2. Configure Policy Action: Add Rule

Policy Name: rd

Rule:

Class Name: tpd

Action: Set CoS (0-7)

Meter

Meter Mode: Rate Limit

Rate (16-1000000): kbps

Apply Revert

FIG. 146 Adding Rules to a Policy Map

Perform these steps 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.

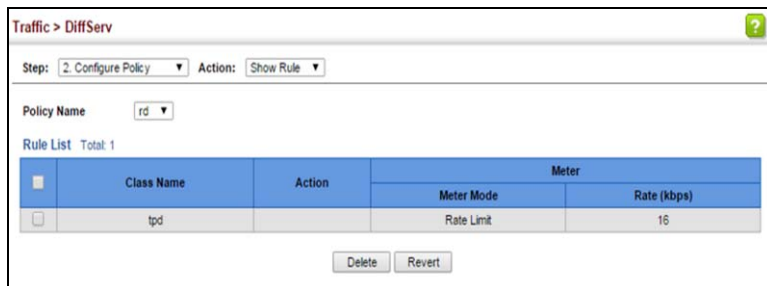


FIG. 147 Showing the Rules for a Policy Map

Attaching a Policy Map to a Port

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

Command Usage

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

The following table lists the options on this page:

Traffic - Diffserv Options	
Port	Specifies a port. (Range: 1-10/26/28/52)
Ingress	Applies the selected rule to ingress traffic.

Perform these steps to bind a policy map to a port:

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

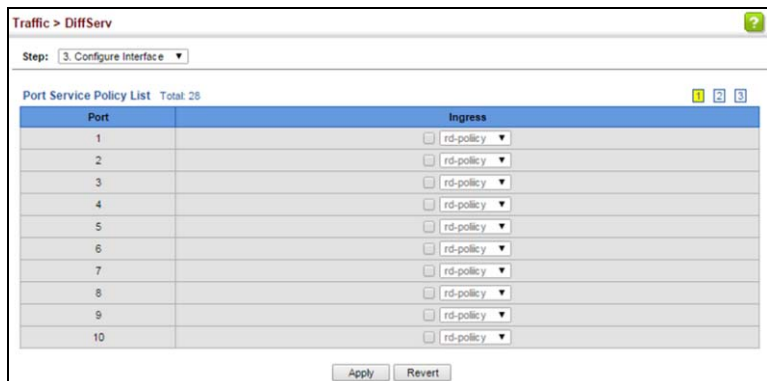


FIG. 148 Attaching a Policy Map to a Port

VoIP Traffic Configuration

This chapter covers the following topics:

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

Overview

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

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

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

Configuring VoIP Traffic

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

Command Usage

All ports are set to VLAN hybrid mode by default. Prior to enabling VoIP for a port (by setting the VoIP mode to Auto or Manual as described below), first ensure that VLAN membership is not set to access mode (see the *Adding Static Members to VLANs* section on page 89).

The following table lists the options on this page:

Traffic - VoIP Options	
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.*

Perform these steps 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**.

The screenshot shows the 'Traffic > VoIP' configuration page. At the top, there is a breadcrumb 'Traffic > VoIP' and a step selector 'Step: 1. Configure Global'. Below this, there are three configuration fields: 'Auto Detection Status' with a checked checkbox and the text 'Enabled'; 'Voice VLAN' with a dropdown menu showing '1234'; and 'Voice VLAN Aging Time (5-43200)' with a text input field containing '3000' and the unit 'sec'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

FIG. 149 Configuring a Voice VLAN

Configuring Telephony OUI

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

The following table lists the options on this page:

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

Perform these steps 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**.

The screenshot shows the 'Traffic > VoIP' configuration page. At the top, 'Step: 2. Configure OUI' and 'Action: Add' are selected. The form contains three fields: 'Telephony OUI' with the value '00-e0-bb-00-00-00', 'Mask' with a dropdown menu showing 'FF-FF-FF-00-00-00', and 'Description' with the text 'old phones'. At the bottom right, there are 'Apply' and 'Revert' buttons.

FIG. 150 Configuring an OUI Telephony List

Perform these steps 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.

The screenshot shows the 'Traffic > VoIP' configuration page with 'Step: 2. Configure OUI' and 'Action: Show' selected. Below the form, a table titled 'Telephony OUI List' shows a total of 3 entries. Each entry has a checkbox, a 'Telephony OUI' value, a 'Mask' value, and a 'Description'.

	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

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

FIG. 151 Showing an OUI Telephony List

Configuring VoIP Traffic Ports

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

Command Usage

All ports are set to VLAN hybrid mode by default. Prior to enabling VoIP for a port (by setting the VoIP mode to Auto or Manual as described below), first ensure that VLAN membership is not set to access mode (see the *Adding Static Members to VLANs* section on page 89).

The following table lists the options on this page:

Traffic - VoIP Options	
Mode	<p>Specifies if the port will be added to the Voice VLAN when VoIP traffic is detected. (Default: None)</p> <ul style="list-style-type: 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	<p>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)</p>
Discovery Protocol	<p>Selects a method to use for detecting VoIP traffic on the port. (Default: OUI)</p> <ul style="list-style-type: none"> 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 the <i>Link Layer Discovery Protocol</i> section on page 190 for more information on LLDP.
Priority	<p>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)</p>
Remaining Age	<p>Number of minutes before this entry is aged out.</p> <p>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.</p> <p>When VoIP Mode is set to Auto, the Remaining Age will be displayed. Otherwise, if the VoIP Mode is Disabled or set to Manual, the remaining age will display NA.</p>

Perform these steps 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**.

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

FIG. 152 Configuring Port Settings for a Voice VLAN

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.
- **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.
- **DHCP Snooping** - Filter IP traffic on insecure ports for which the source address cannot be identified via DHCP snooping.
- **DoS Protection** - Protects against Denial-of-Service attacks.
- **IPv4 Source Guard** - Filters IPv4 traffic on insecure ports for which the source address cannot be identified via DHCPv4 snooping nor static source bindings.

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

AAA (Authentication, Authorization, and Accounting)

The authentication, authorization, and accounting (AAA) feature provides the main framework for configuring access control on the switch. The three security functions can be summarized as follows:

- **Authentication** - Identifies users that request access to the network.
- **Authorization** - Determines if users can access specific services.
- **Accounting** - Provides reports, auditing, and billing for services that users have accessed on the network.

The AAA functions require the use of configured RADIUS or TACACS+ servers in the network. The security servers can be defined as sequential groups that are applied as a method for controlling user access to specified services. For example, when the switch attempts to authenticate a user, a request is sent to the first server in the defined group, if there is no response the second server will be tried, and so on. If at any point a pass or fail is returned, the process stops.

The switch supports the following AAA features:

- Accounting for IEEE 802.1x authenticated users that access the network through the switch.
- Accounting for users that access management interfaces on the switch through the console and Telnet.
- Accounting for commands that users enter at specific CLI privilege levels.
- Authorization of users that access management interfaces on the switch through the console and Telnet.

To configure AAA on the switch, you need to follow this general process:

1. Configure RADIUS and TACACS+ server access parameters. See the *Configuring Local/ Remote Logon Authentication* section on page 135 for more information.
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.

The following table lists the options on this page:

Security - AAA (System Authentication) Options	
Authentication Sequence	Select the authentication, or authentication sequence required: <ul style="list-style-type: none"> • 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.

Perform these steps 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**.



FIG. 153 Configuring the Authentication Sequence

Configuring Remote Login 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 login 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.

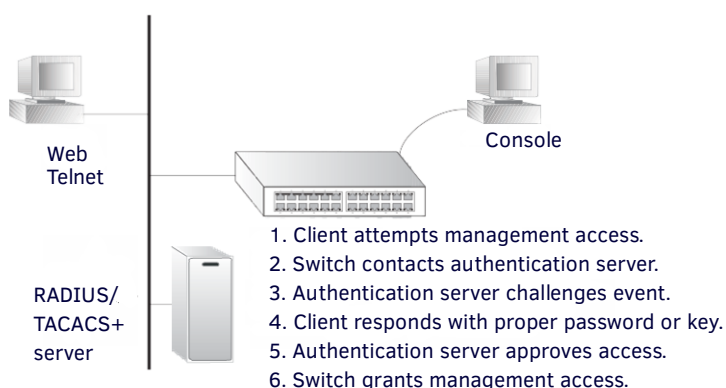


FIG. 154 Authentication Server Operation

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 login authentication control management access via the console port, web browser, or Telnet.
- RADIUS and TACACS+ login 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 login 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).

The following table lists the options on this page:

Security - AAA (Server) Options	
Configure Server	
RADIUS	<ul style="list-style-type: none"> • Global - Provides globally applicable RADIUS settings. • Server Index - Specifies one of five RADIUS servers that may be configured. The switch attempts authentication using the listed sequence of servers. The process ends when a server either approves or denies access to a user. • Server IP Address - Address of authentication server. (A Server Index entry must be selected to display this item.) • Accounting Server UDP Port - Network (UDP) port on authentication server used for accounting messages. (Range: 1-65535; Default: 1813) • Authentication Server UDP Port - Network (UDP) port on authentication server used for authentication messages. (Range: 1-65535; Default: 1812) • Authentication Timeout - The number of seconds the switch waits for a reply from the RADIUS server before it resends the request. (Range: 1-65535; Default: 5) • Authentication Retries - Number of times the switch tries to authenticate logon access via the authentication server. (Range: 1-30; Default: 2) • Set Key - Mark this box to set or modify the encryption key. • Authentication Key - Encryption key used to authenticate logon access for client. Enclose any string containing blank spaces in double quotes. (Maximum length: 48 characters) • Confirm Authentication Key - Re-type the string entered in the previous field to ensure no errors were made. The switch will not change the encryption key if these two fields do not match.
TACACS+	<ul style="list-style-type: none"> • Global - Provides globally applicable TACACS+ settings. • Server Index - Specifies the index number of the server to be configured. The switch currently supports only one TACACS+ server. • Server IP Address - Address of the TACACS+ server. (A Server Index entry must be selected to display this item.) • Authentication Server TCP Port - Network (TCP) port of TACACS+ server used for authentication messages. (Range: 1-65535; Default: 49) • Authentication Timeout - The number of seconds the switch waits for a reply from the TACACS+ server before it resends the request. (Range: 1-65535; Default: 5) • Authentication Retries - Number of times the switch tries to authenticate logon access via the authentication server. (Range: 1-30; Default: 2) • Set Key - Mark this box to set or modify the encryption key. • Authentication Key - Encryption key used to authenticate logon access for client. Enclose any string containing blank spaces in double quotes. (Maximum length: 48 characters) • Confirm Authentication Key - Re-type the string entered in the previous field to ensure no errors were made. The switch will not change the encryption key if these two fields do not match.
Configure Group	
Server Type	Select RADIUS or TACACS+ server.
Group Name	Defines a name for the RADIUS or TACACS+ server group. (Range: 1-64 characters)
Sequence at Priority	Specifies the server and sequence to use for the group. (Range: 1-5 for RADIUS; 1 for TACACS) When specifying the priority sequence for a sever, the server index must already be defined (see the <i>Configuring Local/ Remote Logon Authentication</i> section on page 135).

Perform these steps 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**.

FIG. 155 Configuring Remote Authentication Server (RADIUS)

FIG. 156 Configuring Remote Authentication Server (TACACS+)

Perform these steps 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**.

FIG. 157 Configuring AAA Server Groups

Perform these steps 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.

Security > AAA > Server		
Step:	2. Configure Group	Action: Show
Server Type <input checked="" type="radio"/> RADIUS <input type="radio"/> TACACS+		
RADIUS Group List Total: 3		
<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

FIG. 158 Showing AAA Server Groups

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.

The following table lists the options on this page:

Security - AAA (Accounting) Options	
Configure Global	
Periodic Update	Specifies the interval at which the local accounting service updates information for all users on the system to the accounting server. (Range: 1-2147483647 minutes)
Configure Method	
Accounting Type	Specifies the service as: <ul style="list-style-type: none"> • 802.1x - Accounting for end users. • Command - Administrative accounting to apply to commands entered at specific CLI privilege levels. • Exec - Administrative accounting for local console, Telnet, or SSH connections.
Privilege Level	The CLI privilege levels (0-15). This parameter only applies to Command accounting.
Method Name	Specifies an accounting method for service requests. The default methods are used for a requested service if no other methods have been defined. (Range: 1-64 characters) Note that the method name is only used to describe the accounting method configured on the specified RADIUS or TACACS+ servers. No information is sent to the servers about the method to use.
Accounting Notice	Records user activity from log-in to log-off point.
Server Group Name	Specifies the accounting server group. (Range: 1-64 characters) The group names <i>radius</i> and <i>tacacs+</i> specifies all configured RADIUS and TACACS+ hosts (see the <i>Configuring Local/ Remote Logon Authentication</i> section on page 135.) 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	<ul style="list-style-type: none"> • Method Name - Specifies a user defined accounting method to apply to an interface. This method must be defined in the Configure Method page. (Range: 1-64 characters)
Command	<ul style="list-style-type: none"> • Privilege Level - The CLI privilege levels (0-15). • Console Method Name - Specifies a user-defined method name to apply to commands entered at the specified CLI privilege level through the console interface. • VTY Method Name - Specifies a user-defined method name to apply to commands entered at the specified CLI privilege level through Telnet or SSH.
Exec	<ul style="list-style-type: none"> • Console Method Name - Specifies a user defined method name to apply to console connections. • VTY Method Name - Specifies a user defined method name to apply to Telnet and SSH connections.
Show Information - Summary	
Accounting Type	Displays the accounting service.
Method Name	Displays the user-defined or default accounting method.

Security - AAA (Accounting) Options	
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.

Perform these steps 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**.

Security > AAA > Accounting

Step: 1. Configure Global

Periodic Update (1-2147483647) 1 min

Apply Revert

FIG. 159 Configuring Global Settings for AAA Accounting

Perform these steps to configure the accounting method applied to various service types and the assigned server group:

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

Security > AAA > Accounting

Step: 2. Configure Method Action: Add

Accounting Type: 802.1X

Method Name: default

Accounting Notice: Start-Stop

Server Group Name: radius

Apply Revert

FIG. 160 Configuring AAA Accounting Methods

Perform these steps 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.

The screenshot shows the 'Security > AAA > Accounting' configuration page. The 'Step' is set to '2. Configure Method' and the 'Action' is 'Show'. Below this, there is a 'Method List' table with 18 total entries. The table has columns for Accounting Type, Method Name, Accounting Notice, and Server Group Name. The first row is for '802.1X' with a 'default' method name, 'Start-Stop' notice, and 'radius' server group. The remaining 17 rows are for 'Command' services (Command 0 through Command 8), all with 'default' method names, 'Start-Stop' notices, and 'tacacs+' server groups. There are 'Delete' and 'Revert' buttons at the bottom of the table.

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

FIG. 161 Showing AAA Accounting Methods

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

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

The screenshot shows the 'Security > AAA > Accounting' configuration page. The 'Step' is set to '3. Configure Service'. The 'Accounting Type' is set to '802.1X'. Below this, there is a 'Port Method List' table with 28 total entries. The table has columns for Port and Method Name. The first five rows show ports 1 through 5, all with 'default' method names. There are three pagination buttons (1, 2, 3) at the top right of the table.

Port	Method Name
1	default
2	default
3	default
4	
5	

FIG. 162 Configuring AAA Accounting Service for 802.1x Service

The screenshot shows the 'Security > AAA > Accounting' configuration page. The 'Step' is set to '3. Configure Service'. The 'Accounting Type' is set to 'Command'. Below this, there is a 'Command Method List' table with 16 total entries. The table has columns for Privilege Level, Console Method Name, and VTY Method Name. The first five rows show privilege levels 0 through 4, all with 'default' method names. The last row shows privilege level 5 with 'command4Method' for the console and 'command5Method' for the VTY. There are two pagination buttons (1, 2) at the top right of the table.

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

FIG. 163 Configuring AAA Accounting Service for Command Service

Security > AAA > Accounting

Step: 3. Configure Service

Accounting Type 802.1X Command EXEC

Console Method Name

VTY Method Name

FIG. 164 Configuring AAA Accounting Service for Exec Service

Perform these steps 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**.

Security > AAA > Accounting

Step: 4. Show Information

Summary Statistics

Method List Total: 18

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

FIG. 165 Displaying a Summary of Applied AAA Accounting Methods

Perform these steps 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**.

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

FIG. 166 Displaying Statistics for AAA Accounting Sessions

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.

The following table lists the options on this page:

Security - AAA (Authorization) Options	
Configure Method	
Authorization Type	Specifies the service as: <ul style="list-style-type: none"> • Command - Administrative authorization to apply to commands entered at specific CLI privilege levels. • Exec - Administrative authorization for local console, Telnet, or SSH connections.
Method Name	Specifies an authorization method for service requests. The default method is used for a requested service if no other methods have been defined. (Range: 1-64 characters)
Server Group Name	Specifies the authorization server group. (Range: 1-64 characters) The group name <i>tacacs+</i> specifies all configured TACACS+ hosts (see the <i>Configuring Local/Remote Logon Authentication</i> section on page 135.) Any other group name refers to a server group configured on the TACACS+ Group Settings page. Authorization is only supported for TACACS+ servers.
Configure Service	
Authorization Type	Specifies the service as Exec, indicating administrative authorization for local console, Telnet, or SSH connections.
Console Method Name	Specifies a user defined method name to apply to console connections.
VTY Method Name	Specifies a user defined method name to apply to Telnet and SSH connections.
Show Information	
Authorization Type	Displays the authorization service.
Method Name	Displays the user-defined or default accounting method.
Server Group Name	Displays the authorization server group.
Interface	Displays the console or Telnet interface to which these rules apply. (This field is null if the authorization method and associated server group has not been assigned to an interface.)

Perform these steps 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**.

The screenshot shows the configuration interface for AAA Authorization. At the top, the breadcrumb is 'Security > AAA > Authorization'. Below it, there's a step indicator 'Step: 1. Configure Method' and an 'Action: Add' dropdown. The main configuration area has three fields: 'Authorization Type' with a dropdown menu set to 'EXEC', 'Method Name' with a text input field containing 'default', and 'Server Group Name' with a radio button selected for 'tacacs+' and an empty text input field below it. At the bottom right, there are 'Apply' and 'Revert' buttons.

FIG. 167 Configuring AAA Authorization Methods

Perform these steps 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.

Security > AAA > Authorization

Step: 1. Configure Method Action: Show

Method List Total: 2

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

Delete Revert

FIG. 168 Showing AAA Authorization Methods

Perform these steps 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**.

Security > AAA > Authorization

Step: 2. Configure Service

Authorization Type: EXEC

Console Method Name: tps-auth

VTY Method Name: tps-auth

Apply Revert

FIG. 169 Configuring AAA Authorization Methods for Exec Service

Perform these steps to display the configured authorization method and assigned server groups for the Exec service type:

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

Security > AAA > Authorization

Step: 3. Show Information

Method List Total: 3

<input type="checkbox"/>	Authorization Type	Method Name	Server Group Name	Interface
<input type="checkbox"/>	EXEC	default	tacacs+	
<input type="checkbox"/>	EXEC	console	tacacs+	Console
<input type="checkbox"/>	EXEC	telnet	tacacs+	Telnet

FIG. 170 Displaying the Applied AAA Authorization Method

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 on-board agent. You should therefore assign a new administrator password as soon as possible, and store it in a safe place.

The following table lists the options on this page:

Security - User Accounts Options	
User Name	The name of the user. (Maximum length: 32 characters; maximum number of users: 16)
Access Level	<p>Specifies command access privileges. (Range: 0-15)</p> <p>Level 0, 8 and 15 are designed for users (guest), managers (network maintenance), and administrators (top-level access). The other levels can be used to configure specialized access profiles.</p> <p>Level 0-7 provide the same default access to a limited number of commands which display the current status of the switch, as well as several database clear and reset functions. These commands are equivalent to those available under Normal Exec command mode in the CLI.</p> <p>Level 8-14 provide the same default access privileges, including additional commands beyond those provided for Levels 0-7 (equivalent to CLI Normal Exec command mode), and a subset of the configuration commands provided for Level 15 (equivalent to CLI Privileged Exec command mode).</p> <p>Level 15 provides full access to all commands.</p> <p>The privilege level associated with any command can be changed using the <i>privilege</i> command described in the CLI Reference Guide.</p> <p>Any privilege level can access all of the commands assigned to lower privilege levels. For example, privilege level 8 can access all commands assigned to privilege levels 7-0 according to default settings, and to any other commands assigned to levels 7-0 using the <i>privilege</i> command described in the CLI Reference Guide.</p>
Password Type	<p>Specifies the following options:</p> <ul style="list-style-type: none"> • 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 boot or when downloading the configuration file from a TFTP or FTP server. There is no need for you to manually configure encrypted passwords.
Password	Specifies the user password. (Range: 0-32 characters, case sensitive)
Confirm Password	Re-type the string entered in the previous field to ensure no errors were made. The switch will not change the password if these two fields do not match.

Perform these steps 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**.

The screenshot shows the 'Security > User Accounts' configuration page. At the top, there is a breadcrumb 'Security > User Accounts' and an 'Action' dropdown menu set to 'Add'. Below this, there are several input fields: 'User Name' with the value 'bob', 'Access Level' with a dropdown menu showing '15 (Privileged)', 'Password Type' with a dropdown menu showing 'Plain Password', 'Password' with a masked input field (*****), and 'Confirm Password' with a masked input field (*****). At the bottom of the form, there are two buttons: 'Apply' and 'Revert'.

FIG. 171 Configuring User Accounts

Perform these steps to show user accounts:

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

Security > User Accounts		
Action: Show		
User Account List Total: 3		
<input type="checkbox"/>	User Name	Access Level
<input type="checkbox"/>	admin	15
<input type="checkbox"/>	guest	0
<input type="checkbox"/>	bob	15

Delete Revert

FIG. 172 Showing User Accounts

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 the *Configuring Remote Login Authentication Servers* section on page 135.)

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:

Dynamic QoS Profiles		
Profile	Attribute Syntax	Example
DiffServ	service-policy-in =policy-map-name	service-policy-in=p1
Rate Limit	rate-limit-input =rate rate-limit-output =rate	rate-limit-input=100 (kbps) rate-limit-output=200 (kbps)
802.1p	switchport-priority-default =value	switchport-priority-default=2
IP ACL	ip-access-group-in =ip-acl-name	ip-access-group-in=ipv4acl
IPv6 ACL	ipv6-access-group-in =ipv6-acl-name	ipv6-access-group-in=ipv6acl
MAC ACL	mac-access-group-in =mac-acl-name	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 *service-policy-in=pp1;rate-limit-input=100* attribute 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 (cannot 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 re-authentication time.

The following table lists the options on this page:

Security - Network Access Options	
Aging Status	<p>Enables aging for authenticated MAC addresses stored in the secure MAC address table. (Default: Disabled)</p> <p>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 174).</p> <p>Authenticated MAC addresses are stored as dynamic entries in the switch's secure MAC address table and are removed when the aging time expires.</p> <p>The maximum number of secure MAC addresses supported for the switch system is 1024.</p>

Perform these steps to configure aging status and re-authentication 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 re-authentication time as required.
4. Click **Apply**.



FIG. 173 Configuring Global Settings for Network Access

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.

The following table lists the options on this page:

Security - Network Access Options	
Guest VLAN	<p>Specifies the VLAN to be assigned to the port when 802.1x Authentication or MAC authentication fails. (Range: 0-4094, where 0 means disabled; Default: Disabled)</p> <p>The VLAN must already be created and active (see the <i>Configuring VLAN Groups</i> section on page 88). Also, when used with 802.1x authentication, intrusion action must be set for Guest VLAN (see the <i>Configuring Port Authenticator Settings for 802.1x</i> section on page 174).</p> <p>A port can only be assigned to the guest VLAN in case of failed authentication, and switchport mode is set to Hybrid. (See the <i>Adding Static Members to VLANs</i> section on page 89.)</p>
Dynamic VLAN	<p>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)</p> <p>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.</p> <p>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.</p> <p>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.</p>
MAC Filter ID	<p>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 in the <i>Configuring a MAC Address Filter</i> section on page 148). (Range: 1-64; Default: None)</p>

Perform these steps 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. Set the guest VLAN to use when MAC Authentication or 802.1x Authentication fails, the dynamic VLAN, and the MAC filter.
5. Click **Apply**.

The screenshot shows the 'Security > Network Access' configuration page. The 'Step' dropdown is set to '2. Configure Interface'. Below this is a 'Port List' table with a total of 28 ports. The table has four columns: 'Port', 'Guest VLAN (0-4094, 0: Disabled)', 'Dynamic VLAN', and 'MAC Filter (1-64)'. The first five rows of the table are visible, showing ports 1 through 5. Each row has a 'Port' number, a 'Guest VLAN' value of 0, a 'Dynamic VLAN' checkbox checked and labeled 'Enabled', and a 'MAC Filter' dropdown menu.

Port	Guest VLAN (0-4094, 0: Disabled)	Dynamic VLAN	MAC Filter (1-64)
1	0	<input checked="" type="checkbox"/> Enabled	<input type="text"/>
2	0	<input checked="" type="checkbox"/> Enabled	<input type="text"/>
3	0	<input checked="" type="checkbox"/> Enabled	<input type="text"/>
4	0	<input checked="" type="checkbox"/> Enabled	<input type="text"/>
5	0	<input checked="" type="checkbox"/> Enabled	<input type="text"/>

FIG. 174 Configuring Interface Settings 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.

The following table lists the options on this page:

Security - Network Access Options	
Filter ID	Adds a filter rule for the specified filter. (Range: 1-64)
MAC Address	The filter rule will check ingress packets against the entered MAC address or range of MAC addresses (as defined by the MAC Address Mask).
MAC Address Mask	The filter rule will check for the range of MAC addresses defined by the MAC bit mask. If you omit the mask, the system will assign the default mask of an exact match. (Range: 000000000000 - FFFFFFFF; Default: FFFFFFFF)

Perform these steps 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**.

Security > Network Access

Step: 3. Configure MAC Filter Action: Add

Filter ID (1-64) 22

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

MAC Address Mask FFFFFFFF

Apply Revert

FIG. 175 Configuring a MAC Address Filter for Network Access

Perform these steps 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.

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

FIG. 176 Showing the MAC Address Filter Table for Network Access

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.

The following table lists the options on this page:

Security - Network Access Options	
Query By	Specifies parameters to use in the MAC address query. <ul style="list-style-type: none"> 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	<ul style="list-style-type: none"> 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.

Perform these steps 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**.

Security > Network Access

Step: 4. Show Information

Query by:

Sort Key: MAC Address

MAC Address

Interface

Attribute

Query

Authenticated MAC Address List Total: 8

<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-E1-0D	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"/>	00-01-80-36-95-08	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 32m 22s	Dynamic
<input type="checkbox"/>	00-01-80-3B-D3-7F	Unit 1 / Port 23	10.2.2.10	2008y 20m 12d 11h 22m 28s	Dynamic
<input type="checkbox"/>	00-01-80-3C-3C-19	Unit 2 / Port 23	10.2.2.10	2008y 20m 12d 11h 15m 19s	Dynamic
<input type="checkbox"/>	00-01-80-3C-3E-B3	Unit 2 / Port 23	10.2.2.10	2008y 20m 12d 11h 17m 40s	Dynamic

Delete Revert

FIG. 177 Showing Addresses Authenticated for Network Access

Configuring HTTPS

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

Configuring Global Settings for HTTPS

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

Command Usage

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

A padlock icon should appear in the status bar for Internet Explorer 9, Mozilla Firefox 39, or Google Chrome 44, or more recent versions.
- The following web browsers and operating systems currently support HTTPS:

HTTPS System Support	
Web Browser	Operating System
Internet Explorer 9.x or later	Windows 7, 8, 10
Mozilla Firefox 39 or later	Windows 7, 8, 10, Linux
Google Chrome 44 or later	Windows 7, 8, 10

- To specify a secure-site certificate, see "Replacing the Default Secure-site Certificate" on page 252.

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

The following table lists the options on this page:

Security - HTTPS Options	
HTTPS Status	Allows you to enable/disable the HTTPS server feature on the switch. (Default: Enabled)
HTTPS Port	Specifies the TCP port number used for HTTPS connection to the switch's web interface. (Default: Port 443)

Perform these steps 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**.

The screenshot shows the configuration interface for HTTPS. At the top, it says "Security > HTTPS". Below that, there is a dropdown menu for "Action" with "Configure Global" selected. Underneath, there are two main settings: "HTTPS Status" with a checked checkbox and the text "Enabled", and "UDP Port (1-65535)" with a text input field containing "443". At the bottom right, there are two buttons: "Apply" and "Revert".

FIG. 178 Configuring HTTPS

Replacing the Default Secure-site Certificate

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

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

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

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

NOTE: The switch must be reset for the new certificate to be activated. To reset the switch, see the *Resetting the System* section on page 58 or type `reload` at the command prompt: `Console#reload`

The following table lists the options on this page:

Security - HTTPS Options	
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.

Perform these steps 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**.

FIG. 179 Downloading the Secure-Site Certificate

Configuring the Secure Shell

The Berkeley-standard includes remote access tools originally designed for Unix systems. Some of these tools have also been implemented for Microsoft Windows and other environments. These tools, including commands such as `rlogin` (remote login), `rsh` (remote shell), and `rcp` (remote copy), are not secure from hostile attacks.

Secure Shell (SSH) includes server/client applications intended as a secure replacement for the older Berkeley remote access tools. SSH can also provide remote management access to this switch as a secure replacement for Telnet. When the client contacts the switch via the SSH protocol, the switch generates a public-key that the client uses along with a local user name and password for access authentication. SSH also encrypts all data transfers passing between the switch and SSH-enabled management station clients, and ensures that data traveling over the network arrives unaltered.

NOTE: You need to install an SSH client on the management station to access the switch for management via the SSH protocol.

NOTE: The switch supports 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 135). 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 the *Importing User Public Keys* section on page 154 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 144.) 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)
 - The client sends its password to the server.
 - The switch compares the client's password to those stored in memory.
 - 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

The client sends its RSA public key to the switch.

- The switch compares the client's public key to those stored in memory.
- 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.
- The client uses its private key to decrypt the challenge string, computes the MD5 checksum, and sends the checksum back to the switch.
- 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

- The client first queries the switch to determine if DSA public key authentication using a preferred algorithm is acceptable.
- If the specified algorithm is supported by the switch, it notifies the client to proceed with the authentication process. Otherwise, it rejects the request.
- The client sends a signature generated using the private key to the switch.
- When the server receives this message, it checks whether the supplied key is acceptable for authentication, and if so, it then checks whether the signature is correct. If both checks succeed, the client is authenticated.

NOTE: *The SSH server supports up to eight client sessions. The maximum number of client sessions includes both current Telnet sessions and SSH sessions.*

NOTE: *The SSH server can be accessed using any configured IPv4 or IPv6 interface address on the switch.*

Configuring the SSH Server

Use the Security > SSH (Configure Global) page to enable the SSH server and configure basic settings for authentication.

NOTE: You must generate DSA and RSA host keys before enabling the SSH server. See the *Generating the Host Key Pair* section on page 153.

The following table lists the options on this page:

Security - SSH Options	
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) <ul style="list-style-type: none"> 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.

Perform these steps 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**.

The screenshot shows the 'Security > SSH' configuration page. At the top, it says 'Step: 1. Configure Global'. Below this, there are several configuration options:

- SSH Server Status:** A checkbox labeled 'Enabled' is checked.
- Version:** The value is '2.0'.
- Authentication Timeout (1-120):** A text input field contains '120' followed by 'sec'.
- Authentication Retries (1-5):** A text input field contains '3'.
- Server-Key Size (512-896):** A text input field contains '768'.

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

FIG. 180 Configuring the SSH Server

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 *Importing User Public Keys* section on page 154.

NOTE: A host key pair must be configured on the switch before you can enable the SSH server. See the *Configuring the SSH Server* section on page 153 for more information.

The following table lists the options on this page:

Security - SSH Options	
Host-Key Type	The key type used to generate the host key pair (i.e., public and private keys). (Range: RSA (Version 1), DSA (Version 2), Both; Default: Both) The SSH server uses RSA or DSA for key exchange when the client first establishes a connection with the switch, and then negotiates with the client to select either DES (56-bit) or 3DES (168-bit) for data encryption. NOTE: The switch uses only RSA Version 1 for SSHv1.5 clients and DSA Version 2 for SSHv2 clients.
Save	Saves the host key from RAM (i.e., volatile memory) to flash memory. Otherwise, the host key pair is stored to RAM by default. Note that you must select this item from the Show page. (Default: Disabled)

Perform these steps to generate the SSH host key pair:

1. Click **Security > SSH**.
2. Select **Configure Host Key** from the Step list.
3. Select **Generate** from the Action list.
4. Select the host-key type from the drop-down box.
5. Click **Apply**.

The screenshot shows the 'Security > SSH' configuration page. The 'Step' is '2. Configure Host Key' and the 'Action' is 'Generate'. The 'Host-Key Type' is set to 'Both'. There are 'Apply' and 'Revert' buttons at the bottom.

FIG. 181 Generating the SSH Host Key Pair

Perform these steps to display or clear the SSH host key pair:

1. Click **Security > SSH**.
2. Select **Configure Host Key** from the Step list.
3. Select **Show** from the Action list.
4. Select the option to save the host key from memory to flash by clicking **Save**, or select the host-key type to clear and click **Clear**.

The screenshot shows the 'Security > SSH' configuration page with the 'Action' set to 'Show'. The 'Public-Key of Host-Key' section is expanded, showing RSA and DSA keys. The RSA key is selected, and its public key is displayed in a text area. There are 'Clear' and 'Save' buttons at the bottom.

FIG. 182 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.

The following table lists the options on this page:

Security - SSH Options	
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 the <i>Configuring User Accounts</i> section on page 144).
User Key Type	The type of public key to upload. <ul style="list-style-type: none"> • RSA: The switch accepts a RSA version 1 encrypted public key. • DSA: The switch accepts a DSA version 2 encrypted public key. <p>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.</p> <p>The switch uses only RSA Version 1 for SSHv1.5 clients and DSA Version 2 for SSHv2 clients.</p>
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.

Perform these steps 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**.

The screenshot shows the 'Security > SSH' configuration interface. At the top, the 'Step' is set to '3. Configure User Key' and the 'Action' is 'Copy'. Below this, there are four fields: 'User Name' (steve), 'User-Key Type' (RSA), 'TFTP Server IP Address' (192.168.0.81), and 'Source File Name' (rsa.pub). At the bottom right, there are 'Apply' and 'Revert' buttons.

FIG. 183 Copying the SSH User's Public Key

Perform these steps 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**.

The screenshot shows the 'Security > SSH' configuration interface. The 'Step' is '3. Configure User Key' and the 'Action' is 'Show'. The 'User Name' is 'admin'. Below this, there is a section titled 'Public-Key of User-Key'. It contains two entries: 'RSA' and 'DSA'. The 'RSA' entry has a checkbox and a text area containing a long string of numbers. The 'DSA' entry has a checkbox and a text area containing a long string of alphanumeric characters. At the bottom right, there is a 'Clear' button.

FIG. 184 Showing the SSH User's Public Key

Access Control Lists

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

Configuring Access Control Lists -

An ACL is a sequential list of permit or deny conditions that apply to IP addresses, MAC addresses, or other more specific criteria. This switch tests ingress or egress packets against the conditions in an ACL one by one. A packet will be accepted as soon as it matches a permit rule, or dropped as soon as it matches a deny rule. If no rules match, the packet is accepted.

Command Usage

The following restrictions apply to ACLs:

- The maximum number of ACLs is 512.
- The maximum number of rules per system is 2048 rules.
- An ACL can have up to 2048 rules. However, due to resource restrictions, the average number of rules bound to the ports should not exceed 20.
- The maximum number of rules that can be bound to the ports is 64 for each of the following list types: MAC ACLs, IP ACLs (including Standard and Extended ACLs), IPv6 Standard ACLs, and IPv6 Extended ACLs.

The maximum number of rules (Access Control Entries, or ACEs) stated above is the worst case scenario. In practice, the switch compresses the ACEs in TCAM (a hardware table used to store ACEs), but the actual maximum number of ACEs possible depends on too many factors to be precisely determined. It depends on the amount of hardware resources reserved at runtime for this purpose.

Auto ACE Compression is a software feature used to compress all the ACEs of an ACL to utilize hardware resources more 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.

- If no matches are found down to the end of the list, the traffic is denied. For this reason, frequently hit entries should be placed at the top of the list. There is an implied deny for traffic that is not explicitly permitted. Also, note that a single-entry ACL with only one deny entry has the effect of denying all traffic. You should therefore use at least one permit statement in an ACL or all traffic will be blocked.

Because the switch stops testing after the first match, the order of the conditions is critical. If no conditions match, the packet will be denied.

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

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

Showing TCAM Utilization

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

Command Usage

Policy control entries (PCEs) are used by various system functions which rely on rule-based searches, including Access Control Lists (ACLs), IP Source Guard filter rules, Quality of Service (QoS) processes, QinQ, MAC-based VLANs, VLAN translation, or traps. For example, when binding an ACL to a port, each rule in an ACL will use two PCEs; and when setting an IP Source Guard filter rule for a port, the system will also use two PCEs.

The following table lists the options on this page:

Security - ACL Options	
Pool Capability Code	Abbreviation for processes shown in the TCAM List
Unit	Stack unit identifier
Device	Memory chip used for indicated pools
Pool	Rule slice (or call group). Each slice has a fixed number of rules that are used for the specified features.
Total	The maximum number of policy control entries allocated to the each pool.

Security - ACL Options	
Used	The number of policy control entries used by the operating system.
Free	The number of policy control entries available for use.
Capability	The processes assigned to each pool.

Perform these steps 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.

Security > ACL

Step: 2. Configure ACL Action: Show TCAM

Pool Capability Code:

AM - MAC ACL, A4 - IPv4 ACL, A6S - IPv6 Standard ACL,
 A6E - IPv6 extended ACL, DM - MAC diffServ, D4 - IPv4 diffServ,
 D6S - IPv6 standard diffServ, D6E - IPv6 extended diffServ,
 I - IP source guard, C - CPU interface, L - Link local,
 Reserved - Reserved, ALL - All supported function,

TCAM List Total: 12

Unit	Device	Pool	Total	Used	Free	Capability
1	0	0	64	0	64	A6S
1	0	1	64	0	64	A6E
1	0	2	128	0	128	A4
1	0	3	128	0	128	AM
1	0	4	64	0	64	D6S D6E
1	0	5	128	0	128	D4
1	0	6	128	0	128	DM
1	0	7	128	128	0	Reserved
1	0	8	64	64	0	I
1	0	9	64	64	0	C
1	0	10	64	64	0	Reserved
1	0	11	64	64	0	L

FIG. 185 Showing TCAM Utilization

Setting the ACL Name and Type

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

The following table lists the options on this page:

Security - ACL Options	
ACL Name	Name of the ACL. (Maximum length: 32 characters)
Type	<p>The following filter modes are supported:</p> <ul style="list-style-type: none"> • 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 the <i>ARP Inspection</i> section on page 167 for more information.)

Perform these steps 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**.

FIG. 186 Creating an ACL

Perform these steps 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.

ACL Name	Type
aciStandard1	IP Standard
aciStandard2	IP Standard
aciExtended1	IP Extended
aciExtended2	IP Extended
aciMAC1	MAC
aciMAC2	MAC
aciMAC2	MAC
aciIPv6Standard	IPv6 Standard
aciIPv6Extended	IPv6 Extended
aciARP	ARP

FIG. 187 Showing a List of ACLs

Configuring a Standard IPv4 ACL

Use the Security > ACL (Configure ACL - Add Rule - IP Standard) page to configure a Standard IPv4 ACL.

The following table lists the options on this page:

Security - ACL Options	
Type	Selects the type of ACLs to show in the Name list.
Name	Shows the names of ACLs matching the selected type.
Action	An ACL can contain any combination of permit or deny rules.
Address Type	Specifies the source IP address. Use "Any" to include all possible addresses, "Host" to specify a specific host address in the Address field, or "IP" to specify a range of addresses with the Address and Subnet Mask fields. (Options: Any, Host, IP; Default: Any)
Source IP Address	Source IP address
Source Subnet Mask	A subnet mask containing four integers from 0 to 255, each separated by a period. The mask uses 1 bits to indicate "match" and 0 bits to indicate "ignore." The mask is bitwise ANDed with the specified source IP address, and compared with the address for each IP packet entering the port(s) to which this ACL has been assigned.
Time Range	Name of a time range

Perform these steps 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). If you select Host, enter a specific address. If you select IP, enter a subnet address and the mask for an address range.
8. Click **Apply**.

FIG. 188 Configuring a Standard IPv4 ACL

Configuring an Extended IPv4 ACL

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

The following table lists the options on this page:

Security - ACL Options	
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 158.)
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	<p>Specifies the protocol type to match as TCP, UDP or Others, where others indicates a specific protocol number (0-255). (Options: TCP, UDP, Others; Default: Others)</p> <p>The following items are listed under TCP:</p> <ul style="list-style-type: none"> • 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) <p>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:</p> <ul style="list-style-type: none"> • 1 (fin) - Finish • 2 (syn) - Synchronize • 4 (rst) - Reset • 8 (psh) - Push • 16 (ack) - Acknowledgement • 32 (urg) - Urgent pointer <p>For example, use the code value and mask below to catch packets with the following flags set:</p> <ul style="list-style-type: none"> • 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

Security - ACL Options	
Service Type	<ul style="list-style-type: none"> • Packet priority settings based on the following criteria: • Precedence - IP precedence level. (Range: 0-7) • DSCP - DSCP priority level. (Range: 0-63)
Time Range	Name of a time range

Perform these steps 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**.

FIG. 189 Configuring an Extended IPv4 ACL

Configuring a Standard IPv6 ACL

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

The following table lists the options on this page:

Security - ACL Options	
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

Perform these steps 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). If you select Host, enter a specific address. If you select IPv6-prefix, enter a subnet address and the prefix length.
8. Click **Apply**.

The screenshot shows the 'Security > ACL' configuration interface. At the top, it indicates 'Step: 2. Configure ACL' and 'Action: Add Rule'. Under 'Type', 'IPv6 Standard' is selected. The 'Name' field contains 'R&D#6S'. The 'Action' dropdown is set to 'Permit'. Under 'Source Address Type', 'Host' is selected, and the 'Source IPv6 Address' field contains '2009:DB9:2229::79'. The 'Source Prefix Length (0-128)' field contains '128'. There is a 'Time-Range' checkbox which is unchecked, and a dropdown menu next to it showing 'R&D'. At the bottom, there are 'Apply' and 'Revert' buttons.

FIG. 190 Configuring a Standard IPv6 ACL

Configuring an Extended IPv6 ACL

Use the Security > ACL (Configure ACL - Add Rule - IPv6 Extended) page to configure an Extended IPv6 ACL.

The following table lists the options on this page:

Security - ACL Options	
Type	Selects the type of ACLs to show in the Name list.
Name	Shows the names of ACLs matching the selected type.
Action	An ACL can contain any combination of permit or deny rules.
Source Address Type	Specifies the source IP address type. Use Any to include all possible addresses, Host to specify a specific host address in the Address field, or IPv6-Prefix to specify a range of addresses. (Options: Any, Host, IPv6-Prefix; Default: Any)
Destination Address Type	Specifies the destination IP address type. Use Any to include all possible addresses, or IPv6-Prefix to specify a range of addresses. (Options: Any, IPv6-Prefix; Default: Any)
Source/Destination IPv6 Address	An IPv6 address or network class. The address must be formatted according to RFC 2373 "IPv6 Addressing Architecture," using 8 colon-separated 16-bit hexadecimal values. One double colon may be used in the address to indicate the appropriate number of zeros required to fill the undefined fields.
Source/Destination Prefix-Length	A decimal value indicating how many contiguous bits (from the left) of the address comprise the prefix; i.e., the network portion of the address. (Range: 0-128 bits for the source prefix; 0-8 bits for the destination prefix)
DSCP	DSCP traffic class. (Range: 0-63)
Source Port	Protocol* source port number. (Range: 0-65535) * - Includes TCP, UDP or other protocol types.
Source Port Bit Mask	Decimal number representing the port bits to match. (Range: 0-65535)
Destination Port	Protocol destination port number. (Range: 0-65535)
Destination Port Bit Mask	Decimal number representing the port bits to match. (Range: 0-65535)

Security - ACL Options	
Next Header	<p>Identifies the type of header immediately following the IPv6 header. (Range: 0-255)</p> <p>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 is 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:</p> <ul style="list-style-type: none"> 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

Perform these steps to add rules to an Extended IPv6 ACL:

1. Click **Security > ACL**.
2. Select **Configure ACL** from the Step list.
3. Select **Add Rule** from the Action list.
4. Select **IPv6 Extended** from the Type list.
5. Select the name of an ACL from the Name list.
6. Specify the action (i.e., Permit or Deny).
7. Select the address type (Any or IPv6-prefix). If you select Host, enter a specific address. If you select IPv6-prefix, enter a subnet address and prefix length.
8. Set any other required criteria, such as DSCP or next header type.
9. Click **Apply**.

The screenshot shows the 'Security > ACL' configuration window. At the top, the 'Step' is '2. Configure ACL' and the 'Action' is 'Add Rule'. The main configuration area includes the following fields:

- Action:** Permit (dropdown)
- Source Address Type:** Any (dropdown)
- Source IPv6 Address:** [Empty text box]
- Source Prefix Length (0-128):** 0 (text box)
- Destination Address Type:** Any (dropdown)
- Destination IPv6 Address:** [Empty text box]
- Destination Prefix Length (0-128):** 0 (text box)
- DSCP (0-63):** [Empty text box]
- Next-Header (0-255):** [Empty text box]
- Source Port (0-65535):** [Empty text box]
- Source Port Bit Mask (0-65535):** [Empty text box]
- Destination Port (0-65535):** [Empty text box]
- Destination Port Bit Mask (0-65535):** [Empty text box]
- Time-Range:** (checkbox) R&D (dropdown)

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

FIG. 191 Configuring an Extended IPv6 ACL

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.

The following table lists the options on this page:

Security - ACL Options	
Type	Selects the type of ACLs to show in the Name list.
Name	Shows the names of ACLs matching the selected type.
Action	An ACL can contain any combination of permit or deny rules.
Source/Destination Address Type	Use Any to include all possible addresses, Host to indicate a specific MAC address, or MAC to specify an address range with the Address and Bit Mask fields. (Options: Any, Host, MAC; Default: Any)
Source/Destination MAC Address	Source or destination MAC address
Source/Destination Bit Mask	Hexadecimal mask for source or destination MAC address.
Packet Format	This attribute includes the following packet types: <ul style="list-style-type: none"> Any - Any Ethernet packet type. Untagged-eth2 - Untagged Ethernet II packets. Untagged-802.3 - Untagged Ethernet 802.3 packets. Tagged-eth2 - Tagged Ethernet II packets. Tagged-802.3 - Tagged Ethernet 802.3 packets.
VID	VLAN ID (Range: 1-4094)
VID Bit Mask	VLAN bit mask (Range: 0-4095)
Ethernet Type	This option can only be used to filter Ethernet II formatted packets. (Range: 0-ffff hex.) A detailed listing of Ethernet protocol types can be found in RFC 1060. A few of the more common types include 0800 (IP), 0806 (ARP), 8137 (IPX).
Ethernet Type Bit Mask	Protocol bit mask (Range: 0-ffff hex)
CoS	CoS value (Range: 0-7, where 7 is the highest priority)
CoS Bit Mask	CoS bit mask (Range: 0-7)
Time Range	Name of a time range

Perform these steps 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). 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.
8. Set any other required criteria, such as VID, Ethernet type, or packet format.

9. Click **Apply**.

The screenshot shows the 'Security > ACL' configuration window. The 'Step' is '2. Configure ACL' and the 'Action' is 'Add Rule'. The 'Type' is set to 'MAC'. The 'Name' is 'mac'. The 'Action' is 'Permit'. The 'Source Address Type' is 'Any', 'Source MAC Address' is '00-00-00-00-00-00', 'Source Bit Mask' is '00-00-00-00-00-00', 'Destination Address Type' is 'Any', 'Destination MAC Address' is '00-00-00-00-00-00', and 'Destination Bit Mask' is '00-00-00-00-00-00'. Other fields like VID, CoS, and Time-Range are empty or set to 'R&D'. 'Apply' and 'Revert' buttons are at the bottom.

FIG. 192 Configuring a MAC ACL

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 the *Configuring Global Settings for ARP Inspection* section on page 167).

The following table lists the options on this page:

Security - ACL Options	
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 158.)
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.

Perform these steps 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). 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.
9. Enable logging if required.
10. Click **Apply**.

FIG. 193 Configuring an ARP ACL

Binding a Port to an Access Control List

After configuring ACLs, use the Security > ACL (Configure Interface - Configure) page to bind the ports that need to filter traffic to the appropriate ACLs.

The following table lists the options on this page:

Security - ACL Options	
Type	Selects the type of ACLs to bind to a port.
Port	Port identifier (Range: 1-10/26/28/52)
ACL	ACL used for ingress packets.
Time Range	Name of a time range.
Counter	Enables counter for ACL statistics.

Perform these steps to bind an ACL to a port:

1. Click **Security > ACL**.
2. Select **Configure Interface** from the Step list.
3. Select **Configure** from the Action list.
4. Select **IP, MAC** or **IPv6** from the Type options.
5. Select a port.
6. Select the name of an ACL from the ACL list.
7. Click **Apply**.

FIG. 194 Binding a Port to an ACL

Showing ACL Hardware Counters

Use the Security > ACL > Configure Interface (Show Hardware Counters) page to show statistics for ACL hardware counters. The following table lists the options on this page:

Security - ACL Options	
Port	Port identifier (Range: 1-10/26/28/52)
Type	Selects the type of ACL.
Direction	Displays statistics for ingress or egress traffic.
Query	Displays statistics for selected criteria.
ACL Name	The ACL bound to this port.
Action	Shows if action is to permit or deny specified packets.
Rules	Shows the rules for the ACL bound to this port.
Time-Range	Name of a time range.
Hit	Shows the number of packets matching this ACL.
Clear Counter	Clears the hit counter for the specified ACL.

Perform these steps to show statistics for ACL hardware counters:

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

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

FIG. 195 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 the *DHCP Snooping Global Configuration* section on page 178 for more information). This database is built by DHCP snooping if it is enabled on globally on the switch and on the required VLANs. ARP Inspection can also validate ARP packets against user- configured ARP access control lists (ACLs) for hosts with statically configured addresses (see the *Configuring an ARP ACL* section on page 164).

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.

The following table lists the options on this page:

Security - ARP Inspection Options	
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) <ul style="list-style-type: none"> • Dst-MAC - Validates the destination MAC address in the Ethernet header against the target MAC address in the body of ARP responses. • IP - Checks the ARP body for invalid and unexpected IP addresses. Sender IP addresses are checked in all ARP requests and responses, while target IP addresses are checked only in ARP responses. • Allow Zeros - Allows sender IP address to be 0.0.0.0. • Src-MAC - Validates the source MAC address in the Ethernet header against the sender MAC address in the ARP body. This check is performed on both ARP requests and responses.
Log Message Number	The maximum number of entries saved in a log message. (Range: 0-256; Default: 5)
Log Interval	The interval at which log messages are sent. (Range: 0-86400 seconds; Default: 1 second)

Perform these steps 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**.

FIG. 196 Configuring Global Settings for ARP Inspection

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 164).
- 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.

The following table lists the options on this page:

Security - ARP Inspection Options	
VLAN	Identifier for configured VLANs
DAI Status	Enables Dynamic ARP Inspection for the selected VLAN. (Default: Disabled)
ACL Name	Allows selection of any configured ARP ACLs. (Default: None)
Static	When an ARP ACL is selected, and static mode also selected, the switch only performs ARP Inspection and bypasses validation against the DHCP Snooping Bindings database. When an ARP ACL is selected, but static mode is not selected, the switch first performs ARP Inspection and then validation against the DHCP Snooping Bindings database. (Default: Disabled)

Perform these steps 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**.




FIG. 197 Configuring VLAN Settings for ARP Inspection

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.

The following table lists the options on this page:

Security - ARP Inspection Options	
Interface	Port or trunk identifier
Trust Status	Configures the port as trusted or untrusted. (Default: Untrusted) By default, all untrusted ports are subject to ARP packet rate limiting, and all trusted ports are exempt from ARP packet rate limiting. Packets arriving on trusted interfaces bypass all ARP Inspection and ARP Inspection Validation checks and will always be forwarded, while those arriving on untrusted interfaces are subject to all configured ARP inspection tests.
Packet Rate Limit	Sets the maximum number of ARP packets that can be processed by CPU per second on trusted or untrusted ports. (Range: 0-2048; Default: 15) Setting the rate limit to 0 means that there is no restriction on the number of ARP packets that can be processed by the CPU. The switch will drop all ARP packets received on a port which exceeds the configured ARP-packets-per-second rate limit.

Perform these steps 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**.

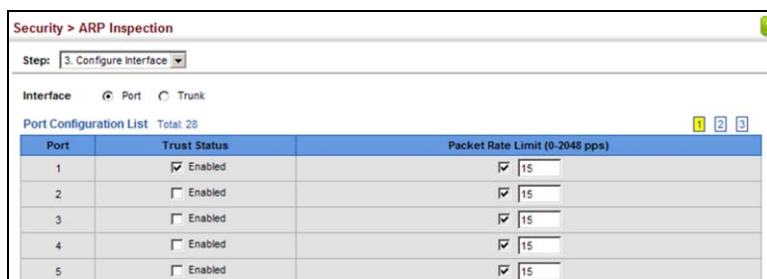


FIG. 198 Configuring Interface Settings for ARP Inspection

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.

The following table lists the available ARP inspection statistics:

ARP Inspection Statistics	
Parameter	Description
Received ARP packets before ARP inspection rate limit	Count of ARP packets received but not exceeding the ARP Inspection rate limit.
Dropped ARP packets in the process of ARP inspection rate limit	Count of ARP packets exceeding (and dropped by) ARP rate limiting.
ARP packets dropped by additional validation (IP)	Count of ARP packets that failed the IP address test.
ARP packets dropped by additional validation (Dst-MAC)	Count of packets that failed the destination MAC address test.
Total ARP packets processed by ARP inspection	Count of all ARP packets processed by the ARP Inspection engine.
ARP packets dropped by additional validation (Src-MAC)	Count of packets that failed the source MAC address test.
ARP packets dropped by ARP ACLs	Count of ARP packets that failed validation against ARP ACL rules.
ARP packets dropped by DHCP snooping	Count of packets that failed validation against the DHCP Snooping Binding database.

Perform these steps 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.

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

FIG. 199 Displaying Statistics for ARP Inspection

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.

The following table lists the available ARP inspection log entries:

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

Perform these steps 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.

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

FIG. 200 Displaying the ARP Inspection Log

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.

The following table lists the options on this page:

Security - IP Filter Options	
Mode	<ul style="list-style-type: none"> • Web - Configures IP address(es) for the web group. • SNMP - Configures IP address(es) for the SNMP group. • Telnet - Configures IP address(es) for the Telnet group. • All - Configures IP address(es) for all groups.
Start IP Address	A single IP address, or the starting address of a range.
End IP Address	The end address of a range.

Perform these steps to create a list of IP addresses authorized for management access:

1. Click **Security > IP Filter**.
2. Select **Add** from the Action list.
3. Select the management interface to filter (Web, SNMP, Telnet, All).
4. Enter the IP addresses or range of addresses that are allowed management access to an interface.
5. Click **Apply**.

FIG. 201 Creating an IP Address Filter for Management Access

Perform these steps to show a list of IP addresses authorized for management access:

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




FIG. 202 Showing IP Addresses Authorized for Management Access

Configuring Port Security

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

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

Command Usage

- The default maximum number of MAC addresses allowed on a secure port is zero (that is, disabled). To use port security, you must configure the maximum number of addresses allowed on a port.
- To configure the maximum number of address entries which can be learned on a port, and then specify the maximum number of dynamic addresses allowed. The switch will learn up to the maximum number of allowed address pairs <source MAC address, VLAN> for frames received on the port. When the port has reached the maximum number of MAC addresses, the port will stop learning new addresses. The MAC addresses already in the address table will be retained and will not be aged out.
Note that you can manually add additional secure addresses to a port using the Static Address Table (page 97).
- 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 61).
- A secure port has the following restrictions:
 - It cannot be used as a member of a static or dynamic trunk.
 - It should not be connected to a network interconnection device.
 - RSPAN and port security are mutually exclusive functions. If port security is enabled on a port, that port cannot be set as an RSPAN uplink port. Also, when a port is configured as an RSPAN uplink port, source port, or destination port, port security cannot be enabled on that port.

The following table lists the options on this page:

Security - Port Security Options	
Port	Port identifier. (Range: 1-10/26/28/52)
Security Status	Enables or disables port security on a port. (Default: Disabled)
Port Status	The operational status: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • None: No action should be taken. (This is the default.) • Trap: Send an SNMP trap message. • Shutdown: Disable the port. • Trap and Shutdown: Send an SNMP trap message and disable the port.
Max MAC Count	The maximum number of MAC addresses that can be learned on a port. (Range: 0 - 1024, where 0 means disabled) The maximum address count is effective when port security is enabled or disabled.
Current MAC Count	The number of MAC addresses currently associated with this interface.
MAC Filter	Shows if MAC address filtering has been set under Security > Network Access (Configure MAC Filter) as described on page 148.

Security - Port Security Options	
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.

Perform these steps 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**.

Port	Security Status	Port Status	Action	Max MAC Count (0-1024)	Current MAC Count	MAC Filter	MAC Filter ID	Last Intrusion MAC	Last Time Detected Intrusion MAC
1	<input checked="" type="checkbox"/> Enabled	Secure/Down	Trap and Shutdown	0	1	Disabled	0	NA	NA
2	<input checked="" type="checkbox"/> Enabled	Secure/Down	Trap	0	0	Disabled	1	NA	NA
3	<input type="checkbox"/> Disabled	Secure/Down	None	0	0	Disabled	0	NA	NA

FIG. 203 Configuring Port Security

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 Message-Digest 5 (MD5), Transport Layer Security (TLS), Protected Extensible Authentication Protocol (PEAP), or Tunneled Transport Layer Security (TTLS). 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 log off message.

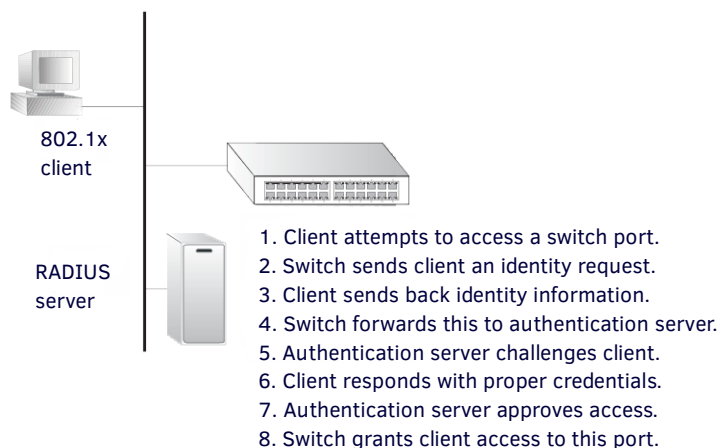


FIG. 204 Configuring Port Authentication

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

- The switch must have an IP address assigned.
- RADIUS authentication must be enabled on the switch and the IP address of the RADIUS server specified.
- 802.1x must be enabled globally for the switch.
- Each switch port that will be used must be set to dot1X Auto mode.
- Each client that needs to be authenticated must have dot1X client software installed and properly configured.
- The RADIUS server and 802.1x client support EAP. (The switch only supports EAPOL in order to pass the EAP packets from the server to the client.)
- The RADIUS server and client also have to support the same EAP authentication type - MD5, PEAP, TLS, or TTLS. (Native support for these encryption methods is provided in Windows 8, 7, Vista and XP, and in Windows 2000 with Service Pack 4. To support these encryption methods in Windows 95 and 98, you can use the AEGIS dot1x client or other comparable client software)

Configuring 802.1x Global Settings

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

The following table lists the options on this page:

Security - Port Authentication Options	
System Authentication Control	Sets the global setting for 802.1x. (Default: Disabled)
Default	Sets all configurable 802.1x global and port settings to their default values.

Perform these steps 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.
4. Click **Apply**.




FIG. 205 Configuring Global Settings for 802.1x Port Authentication

Configuring Port Authenticator Settings for 802.1x

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

Command Usage

- When the switch functions as a local authenticator between supplicant devices attached to the switch and the authentication server, configure the parameters for the exchange of EAP messages between the authenticator and clients on the Authenticator configuration page.
- This switch can be configured to serve as the authenticator on selected ports by setting the Control Mode to Auto on this configuration page, and as a supplicant on other ports by the setting the control mode to Force-Authorized on this page and enabling the PAE supplicant on the Supplicant configuration page.

The following table lists the options on this page:

Security - Port Authentication Options	
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. <ul style="list-style-type: none"> • Yes - Connected client is authorized. • N/A - Connected client is not authorized, or port is not connected.

Security - Port Authentication Options	
Control Mode	<p>Sets the authentication mode to one of the following options:</p> <ul style="list-style-type: none"> • 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	<p>Allows single or multiple hosts (clients) to connect to an 802.1x-authorized port. (Default: Single-Host)</p> <ul style="list-style-type: none"> • Single-Host - Allows only a single host to connect to this port. • Multi-Host - Allows multiple host to connect to this port. <p>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 log off message.</p> <ul style="list-style-type: none"> • MAC-Based - Allows multiple hosts to connect to this port, with each host needing to be authenticated. <p>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).</p>
Max Count	The maximum number of hosts that can connect to a port when the Multi-Host operation mode is selected. (Range: 1-1024; Default: 5)
Max Request	Sets the maximum number of times the switch port will retransmit an EAP request packet to the client before it times out the authentication session. (Range: 1-10; Default 2)
Quiet Period	Sets the time that a switch port waits after the Max Request Count has been exceeded before attempting to acquire a new client. (Range: 1-65535 seconds; Default: 60 seconds)
Tx Period	Sets the time period during an authentication session that the switch waits before re-transmitting an EAP packet. (Range: 1-65535; Default: 30 seconds)
Supplicant Timeout	<p>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)</p> <p>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 re-authentication.</p>
Server Timeout	<p>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)</p> <p>A RADIUS server must be set before the correct operational value of 10 seconds will be displayed in this field. (See the <i>Configuring Remote Login Authentication Servers</i> section on page 135 for more information.)</p>
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	<p>Sets the port's response to a failed authentication.</p> <ul style="list-style-type: none"> • 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 the <i>Configuring VLAN Groups</i> section on page 88) and mapped on each port (see the <i>Configuring Network Access for Ports</i> section on page 147).
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)

Security - Port Authentication Options	
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)

Perform these steps to configure port authenticator settings for 802.1x:

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

Supplicant	Authenticator PAE State Machine	Backend State Machine	Reauthentication State Machine				
State	Reauth Count	Current Identifier	State				
State	Request Count	Identifier (Server)	State				
00-00-00-00-00-00	initialize	0	0	initialize	0	0	initialize

FIG. 206 Configuring Interface Settings for 802.1x Port Authenticator

Displaying 802.1x Statistics

Use the Security > Port Authentication (Show Statistics) page to display statistics for dot1x protocol exchanges for any port. The following table lists the available 802.1x statistics:

802.1x Statistics	
Parameter	Description
<i>Authenticator</i>	
Rx EAPOL Start	The number of EAPOL Start frames that have been received by this Authenticator.
Rx EAPOL Logoff	The number of EAPOL Logoff frames that have been received by this Authenticator.
Rx EAPOL Invalid	The number of EAPOL frames that have been received by this Authenticator in which the frame type is not recognized.
Rx EAPOL Total	The number of valid EAPOL frames of any type that have been received by this Authenticator.
Rx Last EAPOLVer	The protocol version number carried in the most recent EAPOL frame received by this Authenticator.
Rx Last EAPOLSrc	The source MAC address carried in the most recent EAPOL frame received by this Authenticator.
Rx EAP Resp/Id	The number of EAP Resp/Id frames that have been received by this Authenticator.
Rx EAP Resp/Oth	The number of valid EAP Response frames (other than Resp/Id frames) that have been received by this Authenticator.
Rx EAP LenError	The number of EAPOL frames that have been received by this Authenticator in which the Packet Body Length field is invalid.
Tx EAP Req/Id	The number of EAP Req/Id frames that have been transmitted by this Authenticator.
Tx EAP Req/Oth	The number of EAP Request frames (other than Req/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.

802.1x Statistics	
<i>Supplicant</i>	
Rx EAPOL Invalid	The number of EAPOL frames that have been received by this Supplicant in which the frame type is not recognized.
Rx EAPOL Total	The number of valid EAPOL frames of any type that have been received by this Supplicant.
Rx Last EAPOLVer	The protocol version number carried in the most recent EAPOL frame received by this Supplicant.
Rx Last EAPOLSrc	The source MAC address carried in the most recent EAPOL frame received by this Supplicant.
Rx EAP Resp/Id	The number of EAP Resp/Id frames that have been received by this Supplicant.
Rx EAP Resp/Oth	The number of valid EAP Response frames (other than Resp/Id frames) that have been received by this Supplicant.
Rx EAP LenError	The number of EAPOL frames that have been received by this Supplicant in which the Packet Body Length field is invalid.
Tx EAPOL Total	The number of EAPOL frames of any type that have been transmitted by this Supplicant.
Tx EAPOL Start	The number of EAPOL Start frames that have been transmitted by this Supplicant.
Tx EAPOL Logoff	The number of EAPOL Logoff frames that have been transmitted by this Supplicant.
Tx EAP Req/Id	The number of EAP Req/Id frames that have been transmitted by this Supplicant.
Tx EAP Req/Oth	The number of EAP Request frames (other than Req/Id frames) that have been transmitted by this Supplicant.

Perform these steps to display port authenticator statistics for 802.1x:

1. Click **Security > Port Authentication**.
2. Select **Show Statistics** from the Step list.

Security > Port Authentication			
Step:	3. Show Statistics		
Port	1		
Port Authentication Authenticator Statistics			
Rx EAPOL Start	0	Rx EAP Resp/Id	0
Rx EAPOL Logoff	0	Rx EAP Resp/Oth	0
Rx EAPOL Invalid	0	Rx EAP LenError	0
Rx EAPOL Total	0	Tx EAP Req/Id	0
Rx Last EAPOLVer	0	Tx EAP Req/Oth	0
Rx Last EAPOLSrc	00-00-00-00-00-00	Tx EAPOL Total	0
Refresh			

FIG. 207 Showing Statistics for 802.1x Port Authenticator

DHCP Snooping

The addresses assigned to DHCP clients on insecure ports can be carefully controlled using the dynamic bindings registered with DHCP Snooping (or using the static bindings configured with IP Source Guard). DHCP snooping allows a switch to protect a network from rogue DHCP servers or other devices which send port-related information to a DHCP server. This information can be useful in tracking an IP address back to a physical port.

Command Usage

DHCP Snooping Process

- Network traffic may be disrupted when malicious DHCP messages are received from an outside source. DHCP snooping is used to filter DHCP messages received on a non-secure interface from outside the network or fire wall. When DHCP snooping is enabled globally and enabled on a VLAN interface, DHCP messages received on an untrusted interface from a device not listed in the DHCP snooping table will be dropped.
- Table entries are only learned for trusted interfaces. An entry is added or removed dynamically to the DHCP snooping table when a client receives or releases an IP address from a DHCP server. Each entry includes a MAC address, IP address, lease time, VLAN identifier, and port identifier.
- The rate limit for the number of DHCP messages that can be processed by the switch is 100 packets per second. Any DHCP packets in excess of this limit are dropped.
- When DHCP snooping is enabled, DHCP messages entering an untrusted interface are filtered based upon dynamic entries learned via DHCP snooping.

Filtering rules are implemented as follows:

- If the global DHCP snooping is disabled, all DHCP packets are forwarded.
- If DHCP snooping is enabled globally, and also enabled on the VLAN where the DHCP packet is received, all DHCP packets are forwarded for a trusted port. If the received packet is a DHCP ACK message, a dynamic DHCP snooping entry is also added to the binding table.
- If DHCP snooping is enabled globally, and also enabled on the VLAN where the DHCP packet is received, but the port is not trusted, it is processed as follows:
 - If the DHCP packet is a reply packet from a DHCP server (including OFFER, ACK or NAK messages), the packet is dropped.
 - If the DHCP packet is from a client, such as a DECLINE or RELEASE message, the switch forwards the packet only if the corresponding entry is found in the binding table.
 - If the DHCP packet is from a client, such as a DISCOVER, REQUEST, INFORM, DECLINE or RELEASE message, the packet is forwarded if MAC address verification is disabled. However, if MAC address verification is enabled, then the packet will only be forwarded if the client's hardware address stored in the DHCP packet is the same as the source MAC address in the Ethernet header.
 - If the DHCP packet is not a recognizable type, it is dropped.
 - If a DHCP packet from a client passes the filtering criteria above, it will only be forwarded to trusted ports in the same VLAN.
 - If a DHCP packet is from server is received on a trusted port, it will be forwarded to both trusted and untrusted ports in the same VLAN.
 - If the DHCP snooping is globally disabled, all dynamic bindings are removed from the binding table.
 - *Additional considerations when the switch itself is a DHCP client* - The port(s) through which the switch submits a client request to the DHCP server must be configured as trusted. Note that the switch will not add a dynamic entry for itself to the binding table when it receives an ACK message from a DHCP server. Also, when the switch sends out DHCP client packets for itself, no filtering takes place. However, when the switch receives any messages from a DHCP server, any packets received from untrusted ports are dropped.

DHCP Snooping Option 82

- DHCP provides a relay mechanism for sending information about its DHCP clients or the relay agent itself to the DHCP server. Also known as DHCP Option 82, it allows compatible DHCP servers to use the information when assigning IP addresses, or to set other services or policies for clients. It is also an effective tool in preventing malicious network attacks from attached clients on DHCP services, such as IP Spoofing, Client Identifier Spoofing, MAC Address Spoofing, and Address Exhaustion.
- DHCP Snooping must be enabled for Option 82 information to be inserted into request packets.
- When the DHCP Snooping Information Option 82 is enabled, the requesting client (or an intermediate relay agent that has used the information fields to describe itself) can be identified in the DHCP request packets forwarded by the switch and in reply packets sent back from the DHCP server. This information may specify the MAC address or IP address of the requesting device (that is, the switch in this context).
- By default, the switch also fills in the Option 82 circuit-id field with information indicating the local interface over which the switch received the DHCP client request, including the port and VLAN ID. This allows DHCP client-server exchange messages to be forwarded between the server and client without having to flood them to the entire VLAN.
- If DHCP Snooping Information Option 82 is enabled on the switch, information may be inserted into a DHCP request packet received over any VLAN (depending on DHCP snooping filtering rules). The information inserted into the relayed packets includes the circuit-id and remote-id, as well as the gateway Internet address.
- When the switch receives DHCP packets from clients that already include DHCP Option 82 information, the switch can be configured to set the action policy for these packets. The switch can either drop the DHCP packets, keep the existing information, or replace it with the switch's relay information.

DHCP Snooping Global Configuration

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

The following table lists the options on this page:

Security - DHCP Snooping Options	
General	
DHCP Snooping Status	Enables DHCP snooping globally. (Default: Disabled)
SHCP 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)
Information	
DHCP Snooping Information Option Status	Enables or disables DHCP Option 82 information relay. (Default: Disabled)
DHCP Snooping Information Option Sub-option Format	Enables or disables use of sub-type and sub-length fields in circuit-ID (CID) and remote-ID (RID) in Option 82 information. (Default: Enabled)

Security - DHCP Snooping Options	
DHCP Snooping Information Option Remote ID	<p>Specifies the MAC address, IP address, or arbitrary identifier of the requesting device (i.e., the switch in this context).</p> <ul style="list-style-type: none"> 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	<p>Specifies how to handle DHCP client request packets which already contain Option 82 information.</p> <ul style="list-style-type: none"> 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.)

Perform these steps to configure global settings for DHCP Snooping:

1. Click **Security > DHCP Snooping**.
2. Select **Configure Global** from the Step list.
3. Select the required options for the general DHCP snooping process and for the DHCP snooping information option.
4. Click **Apply**.

FIG. 208 Configuring Global Settings for DHCP Snooping

DHCP Snooping VLAN Configuration

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

Command Usage

- When DHCP snooping is enabled globally on the switch, and enabled on the specified VLAN, DHCP packet filtering will be performed on any untrusted ports within the VLAN.
- When the DHCP snooping is globally disabled, DHCP snooping can still be configured for specific VLANs, but the changes will not take effect until DHCP snooping is globally re-enabled.
- When DHCP snooping is globally enabled, and DHCP snooping is then disabled on a VLAN, all dynamic bindings learned for this VLAN are removed from the binding table.

The following table lists the options on this page:

Security - DHCP Snooping Options	
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)

Perform these steps to configure global settings for DHCP Snooping:

1. Click **Security > DHCP Snooping**.
2. Select **Configure VLAN** from the Step list.
3. Enable DHCP Snooping on any existing VLAN.
4. Click **Apply**.

FIG. 209 Configuring DHCP Snooping on a VLAN

Configuring Ports for DHCP Snooping

Use the Security > DHCP Snooping (Configure Interface) page to configure switch ports as trusted or untrusted.

Command Usage

- A trusted interface is an interface that is configured to receive only messages from within the network. An untrusted interface is an interface that is configured to receive messages from outside the network or fire wall.
- When DHCP snooping is enabled both globally and on a VLAN, DHCP packet filtering will be performed on any untrusted ports within the VLAN.
- When an untrusted port is changed to a trusted port, all the dynamic DHCP snooping bindings associated with this port are removed.
- Set all ports connected to DHCP servers within the local network or fire wall to trusted state. Set all other ports outside the local network or fire wall to untrusted state.

The following table lists the options on this page:

Security - DHCP Snooping Options	
Trust Status	Enables or disables a port as trusted. (Default: Disabled)
Max Number	The maximum number of DHCP clients which can be supported per interface. (Range: 1-32; Default: 16)
Circuit ID	Specifies DHCP Option 82 circuit ID sub-option information. <ul style="list-style-type: none"> • Mode - Specifies the default string <i>VLAN-Unit-Port</i> or an arbitrary string. (Default: VLAN-Unit-Port) • Value - An arbitrary string inserted into the circuit identifier field. (Range: 1-32 characters)

Perform these steps to configure global settings for DHCP Snooping:

1. Click **Security > DHCP Snooping**.
2. Select **Configure Interface** from the Step list.
3. Display the list of ports or trunks.
4. Configure the trust status, maximum number of supported clients, and the circuit identifier.
5. Click **Apply**.

FIG. 210 Configuring the Port Mode for DHCP Snooping

Displaying DHCP Snooping Binding Information

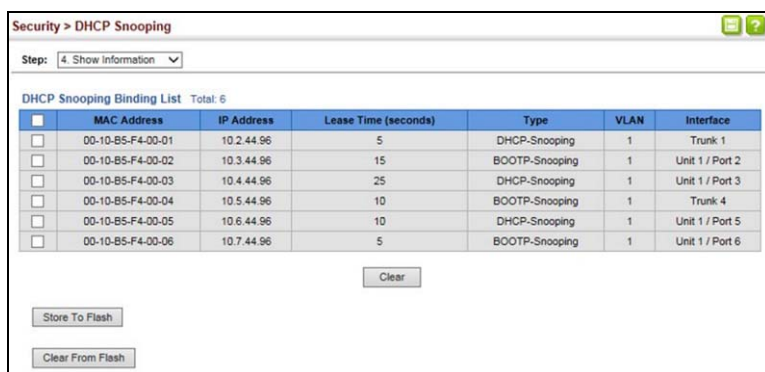
Use the Security > DHCP Snooping (Show Information) page to display entries in the binding table.

The following table lists the options on this page:

Security - DHCP Snooping Options	
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: <ul style="list-style-type: none"> DHCP-Snooping - Dynamically snooped Static-DHCPSNP - Statically configured
VLAN	VLAN to which this entry is bound.
Interface	Port or trunk to which this entry is bound.
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.

Perform these steps to display the binding table for DHCP Snooping:

1. Click **Security > DHCP Snooping**.
2. Select **Show Information** from the Step list.
3. Use the **Store** or **Clear** function, if required.



The screenshot shows the 'Security > DHCP Snooping' configuration page. At the top, there is a 'Step:' dropdown menu set to '4 Show Information'. Below this is a table titled 'DHCP Snooping Binding List' with a total of 6 entries. The table has columns for MAC Address, IP Address, Lease Time (seconds), Type, VLAN, and Interface. Below the table are three buttons: 'Clear', 'Store To Flash', and 'Clear From Flash'.

	MAC Address	IP Address	Lease Time (seconds)	Type	VLAN	Interface
<input type="checkbox"/>	00-10-B5-F4-00-01	10.2.44.96	5	DHCP-Snooping	1	Trunk 1
<input type="checkbox"/>	00-10-B5-F4-00-02	10.3.44.96	15	BOOTP-Snooping	1	Unit 1 / Port 2
<input type="checkbox"/>	00-10-B5-F4-00-03	10.4.44.96	25	DHCP-Snooping	1	Unit 1 / Port 3
<input type="checkbox"/>	00-10-B5-F4-00-04	10.5.44.96	10	BOOTP-Snooping	1	Trunk 4
<input type="checkbox"/>	00-10-B5-F4-00-05	10.6.44.96	10	DHCP-Snooping	1	Unit 1 / Port 5
<input type="checkbox"/>	00-10-B5-F4-00-06	10.7.44.96	5	BOOTP-Snooping	1	Unit 1 / Port 6

FIG. 211 Displaying the Binding Table for DHCP Snooping

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.

The following table lists the options on this page:

Security - DoS Protection Options	
Smurf Attack	Attacks in which a perpetrator generates a large amount of spoofed ICMP Echo Request traffic to the broadcast destination IP address (255.255.255.255), all of which uses a spoofed source address of the intended victim. The victim should crash due to the many interrupts required to send ICMP Echo response packets. (Default: Disabled)
TCP Null Scan	A TCP NULL scan message is used to identify listening TCP ports. The scan uses a series of strangely configured TCP packets which contain a sequence number of 0 and no flags. If the target's TCP port is closed, the target replies with a TCP RST (reset) packet. If the target TCP port is open, it simply discards the TCP NULL scan. (Default: Disabled)
TCP-SYN/FINScan	A TCP SYN/FIN scan message is used to identify listening TCP ports. The scan uses a series of strangely configured TCP packets which contain SYN (synchronize) and FIN (finish) flags. If the target's TCP port is closed, the target replies with a TCP RST (reset) packet. If the target TCP port is open, it simply discards the TCP SYN FIN scan. (Default: Disabled)
TCPXmasScan	A so-called TCP XMAS scan message is used to identify listening TCP ports. This scan uses a series of strangely configured TCP packets which contain a sequence number of 0 and the URG, PSH and FIN flags. If the target's TCP port is closed, the target replies with a TCP RST packet. If the target TCP port is open, it simply discards the TCP XMAS scan. (Default: Disabled)

Perform these steps to protect against DoS attacks:

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




FIG. 212 Protecting Against DoS Attacks

IPv4 Source Guard

IPv4 Source Guard is a security feature that filters IP traffic on network interfaces based on manually configured entries in the IP Source Guard table, or dynamic entries in the DHCP Snooping table when enabled (see the *DHCP Snooping* section on page 177). IP source guard can be used to prevent traffic attacks caused when a host tries to use the IPv4 address of a neighbor to access the network. This section describes how to configure IPv4 Source Guard.

Configuring Ports for IPv4 Source Guard

Use the Security > IP Source Guard > General page to set the filtering type based on source IP address, or source IP address and MAC address pairs. It also specifies lookup within the ACL binding table or the MAC address binding table, as well as the maximum number of allowed binding entries for the lookup tables.

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

Command Usage

Filter Type

- Setting source guard mode to SIP (Source IP) or SIP-MAC (Source IP and MAC) enables this function on the selected port. Use the SIP option to check the VLAN ID, source IP address, and port number against all entries in the binding table. Use the SIP-MAC option to check these same parameters, plus the source MAC address. If no matching entry is found, the packet is dropped.

NOTE: *Multicast addresses cannot be used by IP Source Guard.*

- When enabled, traffic is filtered based upon dynamic entries learned via DHCP snooping (see the *DHCP Snooping* section on page 177), or static addresses configured in the source guard binding table.
- If IP source guard is enabled, an inbound packet's IP address (SIP option) or both its IP address and corresponding MAC address (SIP-MAC option) will be checked against the binding table. If no matching entry is found, the packet will be dropped.
- An entry with same MAC address and a different VLAN ID cannot be added to the binding table.
- Filtering rules are implemented as follows:
 - If DHCP snooping is disabled (see page 178), IP source guard will check the VLAN ID, source IP address, port number, and source MAC address (for the SIP-MAC option). If a matching entry is found in the binding table and the entry type is static IP source guard binding, the packet will be forwarded.
 - If DHCP snooping is enabled, IP source guard will check the VLAN ID, source IP address, port number, and source MAC address (for the SIP-MAC option). If a matching entry is found in the binding table and the entry type is static IP source guard binding, or dynamic DHCP snooping binding, the packet will be forwarded.
 - If IP source guard is enabled on an interface for which IP source bindings have not yet been configured (neither by static configuration in the IP source guard binding table nor dynamically learned from DHCP snooping), the switch will drop all IP traffic on that port, except for DHCP packets allowed by DHCP snooping.

The following table lists the options on this page:

Security - IP Source Guard (General) Options	
Filter Type	Configures the switch to filter inbound traffic based source IP address, or source IP address and corresponding MAC address. (Default: None) <ul style="list-style-type: none"> Disabled - Disables IP source guard filtering on the port. SIP - Enables traffic filtering based on IP addresses stored in the binding table. SIP-MAC - Enables traffic filtering based on IP addresses and corresponding MAC addresses stored in the binding table.
Filter Table	Sets the source guard learning model to search for addresses in the ACL binding table or the MAC address binding table. (Default: ACL binding table)
Max Binding Entry	The maximum number of entries that can be bound to an interface. (ACL Table: 1-5, Default: 5; MAC Table: 1-32, Default: 16) This parameter sets the maximum number of address entries that can be mapped to an interface in the binding table, including both dynamic entries discovered by DHCP snooping (see the <i>DHCP Snooping</i> section on page 177) and static entries set by IP source guard (see the <i>Configuring Static Bindings for IPv4 Source Guard</i> section on page 184).

Perform these steps to set the IP Source Guard filter for ports:

1. Click **Security > IP Source Guard > General**.
2. Set the required filtering type, set the table type to use ACL or MAC address binding, and then set the maximum binding entries for each port.
3. Click **Apply**.

Port	Filter Type	Filter Table	ACL Table Max Binding Entry (1-5)	MAC Table Max Binding Entry (1-32)
1	DISABLED	ACL	5	16
2	DISABLED	ACL	5	16
3	DISABLED	ACL	5	16
4	DISABLED	ACL	5	16
5	DISABLED	ACL	5	16
6	DISABLED	ACL	5	16

FIG. 213 Setting the Filter Type for IPv4 Source Guard

Configuring Static Bindings for IPv4 Source Guard

Use the Security > IP Source Guard > Static Binding (Configure ACL Table and Configure MAC Table) pages to bind a static address to a port. Table entries include a MAC address, IP address, lease time, entry type (Static, Dynamic), VLAN identifier, and port identifier. All static entries are configured with an infinite lease time, which is indicated with a value of zero in the table.

Command Usage

- Table entries include a MAC address, IP address, lease time, entry type (Static-IP- SG-Binding, Dynamic-DHCP-Binding), VLAN identifier, and port identifier.
- Static addresses entered in the source guard binding table are automatically configured with an infinite lease time.
- When source guard is enabled, traffic is filtered based upon dynamic entries learned via DHCP snooping, or static addresses configured in the source guard binding table.
- An entry with same MAC address and a different VLAN ID cannot be added to the binding table.

Static bindings are processed as follows:

- A valid static IP source guard entry will be added to the binding table in ACL mode if one of the following conditions is true:
 - If there is no entry with the same VLAN ID and MAC address, a new entry is added to the binding table using the type *static IP source guard binding*.
 - If there is an entry with the same VLAN ID and MAC address, and the type of entry is static IP source guard binding, then the new entry will replace the old one.
 - If there is an entry with the same VLAN ID and MAC address, and the type of the entry is dynamic DHCP snooping binding, then the new entry will replace the old one and the entry type will be changed to static IP source guard binding.
- A valid static IP source guard entry will be added to the binding table in MAC mode if one of the following conditions are true:
 - If there is no binding entry with the same IP address and MAC address, a new entry will be added to the binding table using the type of static IP source guard binding entry.
 - If there is a binding entry with same IP address and MAC address, then the new entry shall replace the old one.
- Only unicast addresses are accepted for static bindings.

The following table lists the options on this page:

Security - IP Source Guard (Static Binding) Options	
Add - Configure ACL Table	
Port	The port to which a static entry is bound.
VLAN	ID of a configured VLAN (Range: 1-4094)
MAC Address	A valid unicast MAC address
IP Address	A valid unicast IP address, including class types A, B or C.
Add - Configure MAC Table	
MAC Address	A valid unicast MAC address
VLAN	ID of a configured VLAN or a range of VLANs. (Range: 1-4094)
IP Address	A valid unicast IP address, including class types A, B or C.
Port	The port to which a static entry is bound. Specify a physical port number or list of port numbers. Separate nonconsecutive port numbers with a comma and no spaces; or use a hyphen to designate a range of port numbers. (Range: 1-10/28)
Show	
MAC Address	Physical address associated with the entry.
IP Address	IP address corresponding to the client.

Security - IP Source Guard (Static Binding) Options	
VLAN	VLAN to which this entry is bound.
Interface	The port to which this entry is bound.

Perform these steps to configure static bindings for IP Source Guard:

1. Click **Security > IP Source Guard > Static Binding**.
2. Select **Configure ACL Table** or **Configure MAC Table** from the Step list.
3. Select **Add** from the Action list.
4. Enter the required bindings for each port.
5. Click **Apply**.

Security > IP Source Guard > Static Binding

Step: 1. Configure ACL Table Action: Add

Port: 1

VLAN: 1

MAC Address: 00-10-b5-f4-d0-01

IP Address: 10.2.44.96

Apply Revert

FIG. 214 Configuring Static Bindings for IPv4 Source Guard

Perform these steps to display static bindings for IP Source Guard:

1. Click **Security > IP Source Guard > Static Binding**.
2. Select **Configure ACL Table** or **Configure MAC Table** from the Step list.
3. Select **Show** from the Action list.

Security > IP Source Guard > Static Binding

Step: 1. Configure ACL Table Action: Add

Port: 1

VLAN: 1

MAC Address: 00-10-b5-f4-d0-01

IP Address: 10.2.44.96

Apply Revert

FIG. 215 Configuring Static Bindings for IPv4 Source Guard

Displaying Information for Dynamic IPv4 Source Guard Bindings

Use the Security > IP Source Guard > Dynamic Binding page to display the source-guard binding table for a selected interface. The following table lists the options on this page:

Security - IP Source Guard (Dynamic Binding) Options	
Query By	
Port	A port on this switch. (Range: 1-10/26/28/52)
VLAN	ID of a configured VLAN (Range: 1-4094)
MAC Address	A valid unicast MAC address
IP Address	A valid unicast IP address, including class 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	IP address corresponding to the client.
Type	Entry types include DHCP-Snooping or BOOTP-Snooping.

Perform these steps 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**.

The screenshot shows the 'Security > IP Source Guard > Dynamic Binding' page. Under 'Query by:', there are four checkboxes: Port, VLAN, MAC Address, and IP Address. The 'Port' and 'VLAN' checkboxes are checked, and their respective dropdown menus are set to '1'. There are also empty input fields for 'MAC Address' and 'IP Address'. A 'Query' button is located below the input fields. Below the query options, the 'Dynamic Binding List' is displayed with a total of 3 entries. The table has columns for VLAN, MAC Address, Interface, IP Address, and Type.

VLAN	MAC Address	Interface	IP Address	Type
1	00-10-B5-F4-00-01	Unit 1 / Port 2	10.2.44.96	DHCP
1	00-10-B5-F4-00-02	Unit 1 / Port 4	10.2.44.97	DHCP
2	00-10-B5-F4-00-03	Unit 1 / Port 7	10.2.44.98	DHCP

FIG. 216 Showing the IPv4 Source Guard Binding Table

Basic Administration Protocols

This chapter describes basic administration tasks including:

- **Event Logging** - Sets conditions for logging event messages to system memory or flash memory, configures conditions for sending trap messages to remote log servers, and configures trap reporting to remote hosts using Simple Mail Transfer Protocol (SMTP).
- **Link Layer Discovery Protocol (LLDP)** - Configures advertisement of basic information about the local switch, or discovery of information about neighboring devices on the local broadcast domain.
- **Power over Ethernet** - Sets the priority and power budget for each port.
- **Simple Network Management Protocol (SNMP)** - Configures switch management through SNMPv1, SNMPv2c or SNMPv3.
- **Remote Monitoring (RMON)** - Configures local collection of detailed statistics or events which can be subsequently retrieved through SNMP.
- **Time Range** - Sets a time range during which various functions are applied, including applied ACLs or PoE
- **Loopback Detection (LBD)** - Detects general loopback conditions caused by hardware problems or faulty protocol settings.

Configuring Event Logging

The switch allows you to control the logging of error messages, including the type of events that are recorded in switch memory, logging to a remote System Log (syslog) server, and displays a list of recent event messages.

System Log Configuration

Use the Administration > Log > System (Configure Global) page to enable or disable event logging, and specify which levels are logged to RAM or flash memory.

Severe error messages that are logged to flash memory are permanently stored in the switch to assist in troubleshooting network problems. Up to 4096 log entries can be stored in the flash memory, with the oldest entries being overwritten first when the available log memory (256 kilobytes) has been exceeded.

The System Logs page allows you to configure and limit system messages that are logged to flash or RAM memory. The default is for event levels 0 to 3 to be logged to flash and levels 0 to 7 to be logged to RAM.

The following table lists the options on this page:

Administration - Log (System) Options	
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) Logging Levels are as follows: <ul style="list-style-type: none"> • 7: Debug - Debugging messages • 6: Informational - Informational messages only • 5: Notice - Normal but significant condition, such as cold start • 4: Warning - Warning conditions (e.g., return false, unexpected return) • 3: Error - Error conditions (e.g., invalid input, default used) • 2: Critical - Critical conditions (e.g., memory allocation, or free memory error - resource exhausted) • 1: Alert - Immediate action needed • 0: Emergency - System unusable NOTE: <i>There are only Level 2, 5 and 6 error messages for the current firmware release.</i>
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: <i>The Flash Level must be equal to or less than the RAM Level.</i> NOTE: <i>All log messages are retained in RAM and Flash after a warm restart (i.e., power is reset through the command interface).</i> NOTE: <i>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).</i>
Command Log Status	Records the commands executed from the CLI, including the execution time and information about the CLI user including the user name, user interface (console port, telnet or SSH), and user IP address. The severity level for this record type is 6 (a number that indicates the facility used by the syslog server to dispatch log messages to an appropriate service.)

Perform these steps 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**.

FIG. 217 Configuring Settings for System Memory Logs

Perform these steps to show the error messages logged to system or flash memory:

1. Click **Administration > Log > System**.
2. Select **Show System Logs** from the Step list.
3. Click **RAM** to display log messages stored in system memory, or **Flash** to display messages stored in flash memory.
This page allows you to scroll through the logged system and event messages. The switch can store up to 2048 log entries in temporary random access memory (RAM; i.e., memory flushed on power reset) and up to 4096 entries in permanent flash memory.

FIG. 218 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.

The following table lists the options on this page:

Administration - Log (Remote) Options	
Remote Log Status	Enables/disables the logging of debug or error messages to the remote logging process. (Default: Disabled)
Logging Facility	Sets the facility type for remote logging of syslog messages. There are eight facility types specified by values of 16 to 23. The facility type is used by the syslog server to dispatch log messages to an appropriate service. The attribute specifies the facility type tag sent in syslog messages (see RFC 3164). This type has no effect on the kind of messages reported by the switch. However, it may be used by the syslog server to process messages, such as sorting or storing messages in the corresponding database. (Range: 16-23, Default: 23)
Logging Trap Level	Limits log messages that are sent to the remote syslog server for all levels up to the specified level. For example, if level 3 is specified, all messages from level 0 to level 3 will be sent to the remote server. (Range: 0-7, Default: 7)
Server IP Address	Specifies the IPv4 or IPv6 address of a remote server which will be sent syslog messages.
Port	Specifies the UDP port number used by the remote server. (Range: 1-65535; Default: 514)

Perform these steps 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**.

FIG. 219 Configuring Settings for Remote Logging of Error Messages

Sending Simple Mail Transfer Protocol Alerts

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

The following table lists the options on this page:

Administration - Log (SMTP) Options	
SMTP Status	Enables/disables the SMTP function. (Default: Enabled)
Severity	Sets the syslog severity threshold level (see table on page 316) used to trigger alert messages. All events at this level or higher will be sent to the configured email recipients. For example, using Level 7 will report all events from level 7 to level 0. (Default: Level 7)
Email Source Address	Sets the email address used for the From field in alert messages. You may use a symbolic email address that identifies the switch, or the address of an administrator responsible for the switch. (Range: 1-41 characters)
Email Destination Address	Specifies the email recipients of alert messages. You can specify up to five recipients.
Server IP Address	Specifies a list of up to three recipient SMTP servers. IPv4 or IPv6 addresses may be specified. The switch attempts to connect to the listed servers in sequential order if the first server fails to respond.

Perform these steps to configure SMTP alert messages:

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

FIG. 220 Configuring SMTP Alert Messages

Link Layer Discovery Protocol

Link Layer Discovery Protocol (LLDP) is used to discover basic information about neighboring devices on the local broadcast domain. LLDP is a Layer 2 protocol that uses periodic broadcasts to advertise information about the sending device. Advertised information is represented in Type Length Value (TLV) format according to the IEEE 802.1AB standard, and can include details such as device identification, capabilities and configuration settings. LLDP also defines how to store and maintain information gathered about the neighboring network nodes it discovers.

Link Layer Discovery Protocol - Media Endpoint Discovery (LLDP-MED) is an extension of LLDP intended for managing endpoint devices such as Voice over IP phones and network switches. The LLDP-MED TLVs advertise information such as network policy, power, inventory, and device location details. LLDP and LLDP-MED information can be used by SNMP applications to simplify troubleshooting, enhance network management, and maintain an accurate network topology.

Setting LLDP Timing Attributes

Use the Administration > LLDP (Configure Global) page to set attributes for general functions such as globally enabling LLDP on the switch, setting the message ageout time, and setting the frequency for broadcasting general advertisements or reports about changes in the LLDP MIB.

The following table lists the options on this page:

Administration - LLDP Options	
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 <code>lldpStatsRemTableLastChangeTime</code> to detect any <code>lldpRemTablesChange</code> 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.

Perform these steps 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**.

Administration > LLDP

Step: 1. Configure Global

LLDP Enabled

Transmission Interval (5-32768) sec

Hold Time Multiplier (2-10)

Delay Interval (1-8192) sec

Reinitialization Delay (1-10) sec

Notification Interval (5-3600) sec

MED Fast Start Count (1-10)

Note: The Transmission Interval must be greater than or equal to 4 times the Delay Interval.

FIG. 221 Configuring LLDP Timing Attributes

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.

The following table lists the options on this page:

Administration - LLDP Options	
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	<p>Enables the transmission of SNMP trap notifications about LLDP and LLDP-MED changes. (Default: Enabled)</p> <p>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.</p> <p>For information on defining SNMP trap destinations, see the <i>Specifying Trap Managers</i> section on page 218.</p> <p>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 <code>lldpStatsRemTableLastChangeTime</code> to detect any <code>lldpRemTablesChange</code> notification-events missed due to throttling or transmission loss.</p>
MED Notification	Enables the transmission of SNMP trap notifications about LLDP-MED changes. (Default: Disabled)

Administration - LLDP Options	
Basic Optional TLVs	<p>Configures basic information included in the TLV field of advertised messages.</p> <ul style="list-style-type: none"> • Management Address - The management address protocol packet includes the IPv4 address of the switch. If no management address is available, the address should be the MAC address for the CPU or for the port sending this advertisement. (Default: Enabled) <p>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.</p> <p>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.</p> <p>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.</p> <ul style="list-style-type: none"> • Port Description - The port description is taken from the ifDescr object in RFC 2863, which includes information about the manufacturer, the product name, and the version of the interface hardware/software. (Default: Enabled) • System Capabilities - The system capabilities identifies the primary function(s) of the system and whether or not these primary functions are enabled. The information advertised by this TLV is described in IEEE 802.1AB. (Default: Enabled) • System Description - The system description is taken from the sysDescr object in RFC 3418, which includes the full name and version identification of the system's hardware type, software operating system, and networking software. (Default: Enabled) • System Name - The system name is taken from the sysName object in RFC 3418, which contains the system's administratively assigned name. To configure the system name, see the <i>System Information</i> section on page 39. (Default: Enabled)
802.1 Organizationally Specific TLVs	<p>Configures IEEE 802.1 information included in the TLV field of advertised messages.</p> <ul style="list-style-type: none"> • Protocol Identity - The protocols that are accessible through this interface (see the <i>Protocol VLANs</i> section on page 92.) (Default: Enabled) • VLAN ID - The port's default VLAN identifier (PVID) indicates the VLAN with which untagged or priority-tagged frames are associated (see the <i>IEEE 802.1Q VLANs</i> section on page 87.) (Default: Enabled) • VLAN Name - The name of all VLANs to which this interface has been assigned (see the <i>IEEE 802.1Q VLANs</i> section on page 87.) (Default: Enabled) • Port and Protocol VLAN ID - The port-based protocol VLANs configured on this interface (see the <i>Protocol VLANs</i> section on page 92.) (Default: Enabled)
802.3 Organizationally Specific TLVs	<p>Configures IEEE 802.3 information included in the TLV field of advertised messages.</p> <ul style="list-style-type: none"> • Link Aggregation - The link aggregation capabilities, aggregation status of the link, and the IEEE 802.3 aggregated port identifier if this interface is currently a link aggregation member. (Default: Enabled) • Max Frame Size - The maximum frame size. (See the <i>Configuring Support for Jumbo Frames</i> section on page 41 for information on configuring the maximum frame size for this switch. (Default: Enabled) • MAC/PHY Configuration/Status - The MAC/PHY configuration and status which includes information about auto-negotiation support/capabilities, and operational Multistation Access Unit (MAU) type. (Default: Enabled) • PoE - Power-over-Ethernet capabilities, including whether or not PoE is supported, currently enabled, if the port pins through which power is delivered can be controlled, the port pins selected to deliver power, and the power class. (Default: Enabled)
MED TLVs	<p>Configures general information included in the MED TLV field of advertised messages.</p> <ul style="list-style-type: none"> • Capabilities - This option advertises LLDP-MED TLV capabilities, allowing Media Endpoint and Connectivity Devices to efficiently discover which LLDP-MED related TLVs are supported on the switch. (Default: Enabled) • Extended Power - This option advertises extended Power-over-Ethernet capability details, such as power availability from the switch, and power state of the switch, including whether the switch is operating from primary or backup power (the Endpoint Device could use this information to decide to enter power conservation mode). (Default: Enabled) • Inventory - This option advertises device details useful for inventory management, such as manufacturer, model, software version and other pertinent information. (Default: Enabled) • Location - This option advertises location identification details. (Default: Enabled) • Network Policy - This option advertises network policy configuration information, aiding in the discovery and diagnosis of VLAN configuration mismatches on a port. Improper network policy configurations frequently result in voice quality degradation or complete service disruption. (Default: Enabled)

Administration - LLDP Options	
MED-Location Civic Address	<p>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.</p> <ul style="list-style-type: none"> 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 can refer to the location of the DHCP server, the location of the network element closest to client, or the location of the client (This is the default.)

Perform these steps 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**.

The screenshot shows the 'Administration > LLDP' configuration page. The 'Step' is '2. Configure Interface' and the 'Action' is 'Configure General'. The interface is set to 'Port 1'. The 'Admin Status' is 'Tx Rx'. 'SNMP Notification' and 'MED Notification' are both disabled. Under 'Basic Optional TLVs', 'Management Address', 'Port Description', 'System Capabilities', 'System Description', and 'System Name' are all disabled. Under '802.1 Organizationally Specific TLVs', 'Protocol Identity', 'VLAN ID', 'VLAN Name', and 'Port and Protocol VLAN ID' are all checked. Under '802.3 Organizationally Specific TLVs', 'Link Aggregation', 'Max Frame Size', and 'MAC/PHY Configuration/Status' are all disabled. 'PoE' is also disabled. Under 'MED TLVs', 'Capabilities', 'Extended Power', 'Inventory', 'Location', and 'Network Policy' are all disabled. The 'MED-Location Civic Address' section shows 'Country' set to 'US' and 'Device entry refers to' set to 'Location of the client'. A note at the bottom states: 'Note: The country string shall be a two-letter ISO 3166 country code, e.g. US'. 'Apply' and 'Revert' buttons are at the bottom.

FIG. 222 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.

LLDP MED Location CA Types		
CA Type	Description	CA Value Example
1	National subdivisions (state, canton, province)	California
2	County, parish	Orange
3	City, township	Irvine
4	City division, borough, city district	West Irvine
5	Neighborhood, block	Riverside
6	Group of streets below the neighborhood level	Exchange
18	Street suffix or type	Avenue
19	House Number	320
20	House number suffix	A
21	Landmark or vanity address	Tech Center
26	Unit (apartment, suite)	Apt519
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.

The following table lists the options on this page:

Administration - LLDP Options	
CA-Type	Descriptor of the data civic address value. (Range: 0-255)
CA-Value	Description of a location. (Range: 1-32 characters)

Perform these steps to specify the physical location of the attached device:

- Click **Administration > LLDP**.
- Select **Configure Interface** from the Step list.
- Select **Add CA-Type** from the Action list.
- Select an interface from the Port or Trunk list.
- Specify a CA-Type and CA-Value pair.
- Click **Apply**.

The screenshot shows the 'Administration > LLDP' configuration page. At the top, it indicates the current step is '2. Configure Interface' and the action is 'Add CA-Type'. Below this, there are radio buttons for 'Port' (selected) and 'Trunk'. Under 'Port', the value '1' is selected. The 'CA-Type (0-255)' field contains the value '1', and the 'CA-Value' field contains the value 'California'. At the bottom of the form, there are two buttons: 'Apply' and 'Revert'.

FIG. 223 Configuring the Civic Address for an LLDP Interface

Perform these step 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.




FIG. 224 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.

The following table lists the options on this page:

Administration - LLDP Options	
General Settings	
Chassis Type	Identifies the chassis containing the IEEE 802 LAN entity associated with the transmitting LLDP agent. There are several ways in which a chassis may be identified and a chassis ID subtype is used to indicate the type of component being referenced by the chassis ID field. The following lists the reference for each chassis ID: <ul style="list-style-type: none"> • 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 the <i>System Information</i> section on page 39).
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. <ul style="list-style-type: none"> • Repeater - IETF RFC 2108 • Bridge - IETF RFC 2674 • WLAN Access Point - IEEE 802.11 MIB • Router - IETF RFC 1812 • Telephone - IETF RFC 2011 • DOCSIS cable device - IETF RFC 2669 and IETF RFC 2670 • End Station Only - IETF RFC 2011
System Capabilities Enabled	The primary function(s) of the system which are currently enabled. Refer to the preceding table.
Management Address	The management address associated with the local system. If no management address is available, the address should be the MAC address for the CPU or for the port sending this advertisement.
Interface Settings	
The attributes listed below apply to both port and trunk interface types. When a trunk is listed, the descriptions apply to the first port of the trunk.	
Port/Trunk Description	A string that indicates the port or trunk description. If RFC 2863 is implemented, the ifDescr object should be used for this field.
Port/Trunk ID	A string that contains the specific identifier for the port or trunk from which this LLDPDU was transmitted.

Administration - LLDP Options	
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	<p>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. The following lists the references for the port ID subtypes:</p> <ul style="list-style-type: none"> • Interface alias - IfAlias (IETF RFC 2863) • Chassis component - EntPhysicalAlias when entPhysClass has a value of 'chassis(3)' (IETF RFC 2737) • Port component - EntPhysicalAlias when entPhysicalClass has a value 'port(10)' or 'backplane(4)' (IETF RFC 2737) • MAC address - MAC address (IEEE Std 802-2001) • Network address - networkAddress • Interface name - ifName (IETF RFC 2863) • Agent circuit ID - agent circuit ID (IETF RFC 3046) • Locally assigned - locally assigned
Port/Trunk ID	A string that contains the specific identifier for the local interface based on interface subtype used by this switch.
Port/Trunk Description	A string that indicates the port or trunk description. If RFC 2863 is implemented, the ifDescr object should be used for this field.
MED Capability	<p>The supported set of capabilities that define the primary function(s) of the interface:</p> <ul style="list-style-type: none"> • LLDP-MED Capabilities • Network Policy • Location Identification • Extended Power via MDI - PSE • Extended Power via MDI - PD • Inventory

Perform these steps to display LLDP information for the local device:

1. Click **Administration > LLDP**.
2. Select **Show Local Device Information** from the Step list.
3. Select **General, Port, Port Details, Trunk, or Trunk Details**.

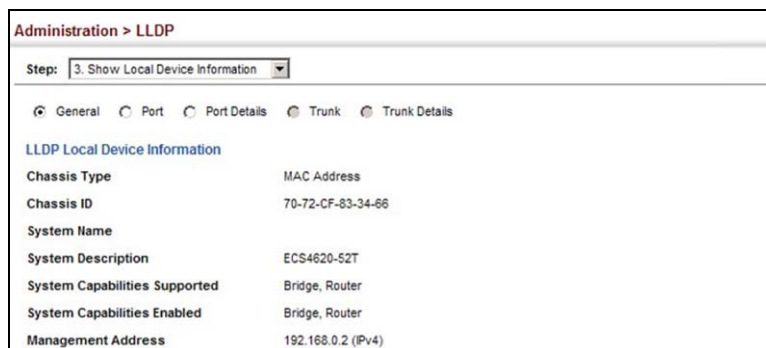


FIG. 225 Displaying Local Device Information for LLDP (General)

The screenshot shows the 'Administration > LLDP' page with the 'Step: 3. Show Local Device Information' dropdown menu. Below the step menu, there are radio buttons for 'General', 'Port', 'Port Details', 'Trunk', and 'Trunk Details', with 'Port' selected. The main content area displays 'LLDP Local Device Port List Total: 28' and a table with the following data:

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

FIG. 226 Displaying Local Device Information for LLDP (Port)

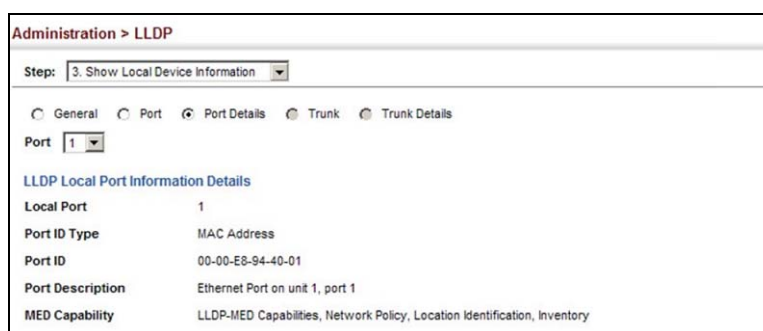


FIG. 227 Displaying Local Device Information for LLDP (Port Details)

Displaying LLDP Remote Device Information

Use the Administration > LLDP (Show Remote Device Information) page to display information about devices connected directly to the switch's ports which are advertising information through LLDP, or to display detailed information about an LLDP-enabled device connected to a specific port on the local switch.

The following table lists the options on this page:

Administration - LLDP Options	
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. (Range: 1-10/26/28/52)
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 <i>Chassis Type</i> on page 195.)
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 <i>Port/Trunk ID Type</i> on page 196.
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 <i>System Capabilities Supported</i> on page 195.)
System Capabilities Enabled	The primary function(s) of the system which are currently enabled. (See <i>System Capabilities Supported</i> on page 195.)
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.

Administration - LLDP Options	
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. <ul style="list-style-type: none"> • 0 - other or unknown • 1 - 10BASE-T half duplex mode • 2 - 10BASE-T full duplex mode • 3 - 100BASE-T4 • 4 - 100BASE-TX half duplex mode • 5 - 100BASE-TX full duplex mode • 6 - 100BASE-T2 half duplex mode • 7 - 100BASE-T2 full duplex mode • 8 - PAUSE for full-duplex links • 9 - Asymmetric PAUSE for full-duplex links • 10 - Symmetric PAUSE for full-duplex links • 11 - Asymmetric and Symmetric PAUSE for full-duplex links • 12 - 1000BASE-X, -LX, -SX, -CX half duplex mode • 13 - 1000BASE-X, -LX, -SX, -CX full duplex mode • 14 - 1000BASE-T half duplex mode • 15 - 1000BASE-T full duplex mode
Remote Port Auto-Neg Status	Shows whether port auto-negotiation is enabled on a port associated with the remote system.
Remote Port MAU Type	An integer value that indicates the operational MAU type of the sending device. This object contains the integer value derived from the list position of the corresponding dot3MauType as listed in IETF RFC 3636 and is equal to the last number in the respective dot3MauType OID.
Port Details - 802.3 Extension Power Information	
Remote Power Class	The port Class of the given port associated with the remote system (PSE - Power Sourcing Equipment or PD - Powered Device).
Remote Power MDI Status	Shows whether MDI power is enabled on the given port associated with the remote system.
Remote Power Pairs	"Signal" means that the signal pairs only are in use, and "Spare" means that the spare pairs only are in use.
Remote Power MDI Supported	Shows whether MDI power is supported on the given port associated with the remote system.
Remote Power Pair Controllable	Indicates whether the pair selection can be controlled for sourcing power on the given port associated with the remote system.
Remote Power Classification	This classification is used to tag different terminals on the Power over LAN network according to their power consumption. Devices such as IP telephones, WLAN access points and others, will be classified according to their power requirements.
Port Details - 802.3 Extension Trunk Information	
Remote Link Aggregation Capable	Shows if the remote port is not in link aggregation state and/or it does not support link aggregation.
Remote Link Aggregation Status	The current aggregation status of the link.
Remote Link Port ID	This object contains the IEEE 802.3 aggregated port identifier, aAggPortID (IEEE 802.3-2002, 30.7.2.1.1), derived from the ifNumber of the ifIndex for the port component associated with the remote system. If the remote port is not in link aggregation state and/or it does not support link aggregation, this value should be zero.
Port Details - 802.3 Extension Frame Information	
Remote Max Frame Size	An integer value indicating the maximum supported frame size in octets on the port component associated with the remote system.

Administration - LLDP Options	
Port Details LLDP-MED Capability	
These fields are only displayed for end-node devices advertising LLDP-MED TLVs.	
Device Class	Any of the following categories of endpoint devices: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • LLDP-MED Capabilities • Network Policy • Location Identification • Extended Power via MDI - PSE • Extended Power via MDI - PD • Inventory
Current Capabilities	The set of capabilities that define the primary function(s) of the port which are currently enabled.
Port Details - Network Policy	
These fields are only displayed for end-node devices advertising LLDP-MED TLVs.	
Application Type	The primary application) defined for this network policy: <ul style="list-style-type: none"> • 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 Identification	
These fields are only displayed for end-node devices advertising LLDP-MED TLVs.	
Location Data Format	Any of these location ID data formats: <ul style="list-style-type: none"> • Coordinate-based LCI - Defined in RFC 3825, includes latitude resolution, latitude, longitude resolution, longitude, altitude type, altitude resolution, altitude, and datum. • Civic Address LCI - Includes What, Country code, CA type, CA length and CA value. "What" is described as the field entry "Device entry refers to" under the <i>Configuring LLDP Interface Attributes</i> section on page 191. The other items and described under the <i>Configuring LLDP Interface Civic-Address</i> section on page 194. • 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 the <i>Configuring LLDP Interface Attributes</i> section on page 191.
Port Details - Extended Power-via-MDI	
Power Type	Power Sourcing Entity (PSE) or Power Device (PD).
Power Priority	Shows power priority for a port. (Unknown, Low, High, Critical)

Administration - LLDP Options	
Power Source	Shows information based on the type of device: <ul style="list-style-type: none"> • PD - Unknown, PSE, Local, PSE and Local • PSE - Unknown, Primary Power Source, Backup Power Source - Power conservation mode
Power Value	The total power in watts required by a PD device from a PSE device, or the total power a PSE device is capable of sourcing over a maximum length cable based on its current configuration. This parameter supports a maximum power required or available value of 102.3 Watts to allow for future expansion. (Range: 0 - 102.3 Watts)
Port Details - Inventory These fields are only displayed for end-node devices advertising LLDP-MED TLVs.	
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.

Perform these steps to display LLDP information for a remote port:

1. Click **Administration > LLDP**.
2. Select **Show Remote Device Information** from the Step list.
3. Select **Port**, **Port Details**, **Trunk**, or **Trunk Details**.
4. When the next page opens, select a port on this switch and the index for a remote device attached to this port.
5. Click **Query**.

The figure consists of two screenshots of the LLDP administration interface. The top screenshot shows the 'Administration > LLDP' page with the step '4. Show Remote Device Information' selected. Below the step list, there are radio buttons for 'Port', 'Port Details', 'Trunk', and 'Trunk Details', with 'Port' selected. A table titled 'LLDP Remote Device Port List' shows one entry with columns: Local Port (1), Chassis ID (70-72-CF-32-0D-FD), Port ID (70-72-CF-32-0D-FF), and System Name (RD93). The bottom screenshot shows the same page with the 'Port' dropdown set to '1' and the 'Remote Index' dropdown set to '2'. A 'Query' button is visible at the bottom.

FIG. 228 Displaying Remote Device Information for LLDP (Port)

Administration > LLDP

Step: 4. Show Remote Device Information

Port
 Port Details
 Trunk
 Trunk Details

Port: 23

Remote Index: 4

Query

LLDP Remote Device Port Information

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

Management Address List Total: 1

Address	Address Type
192.168.0.3	IPv4 Address

802.1 Extension Information

Remote Port VID: 1

Remote Port-Protocol VLAN List Total: 1

VLAN	Support	Status
1	Yes	Enabled

Remote VLAN Name List Total: 1

VLAN	Name
1	DefaultVlan

Remote Protocol Identity List Total: 1

Remote Protocol Identity (Hex)
88-CC

802.3 Extension Port Information

Remote Port Auto-Reg Supported	Yes	Remote Port Auto-Reg Status	Enabled
Remote Port Auto-Reg 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
-----------------------	------

FIG. 229 Displaying Remote Device Information for LLDP (Port Details)

Additional information displayed by an end-point device which advertises LLDP- MED TLVs is shown in FIG. 230.

Administration > LLDP			
Step: 4. Show Remote Device Information			
LLDP-MED Capability			
Device Class	Network Connectivity		
Supported Capabilities	LLDP-MED Capabilities, Network Policy, Location Identification, Inventory		
Current Capabilities	LLDP-MED Capabilities, Network Policy, Location Identification, Inventory		
Network Policy			
Application Type	Guest Voice Signaling	Unknown Policy Flag	Disabled
Tagged Flag	Disabled	VLAN ID	7
Layer 2 Priority	2	DSCP Value	62
Location Identification			
Location Data Format	Coordinate-based LCI		
Country Code	TW	What	2
Location Identification			
Location Data Format	Civic Address LCI		
Country Code	US	What	2
Extended Power-via-MDI			
Power Type	PSE	Power Source	Unknown
Power Priority	Unknown	Power Value	0 W Watts
Inventory			
Hardware Revision	R01	Firmware Revision	1.2.2.1
Software Revision	1.2.2.1	Serial Number	LN10230092
Manufacture Name		Model Name	L
Asset ID			

FIG. 230 Displaying Remote Device Information for LLDP (End Node)

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.

The following table lists the options on this page:

Administration - LLDP Options	
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.

Perform these steps 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**.

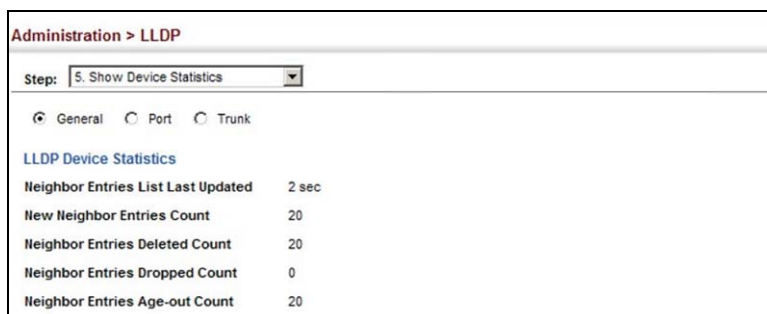


FIG. 231 Displaying LLDP Device Statistics (General)

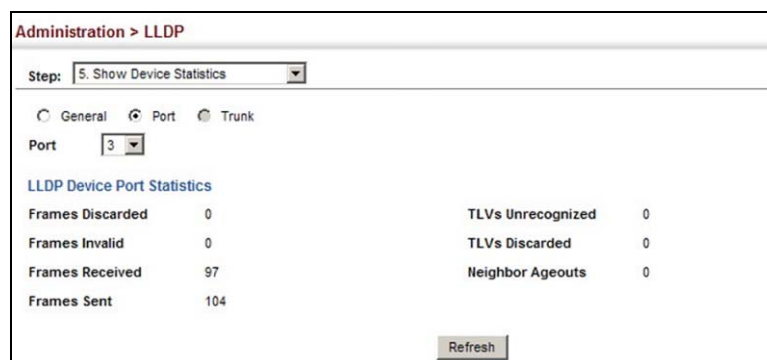


FIG. 232 Displaying LLDP Device Statistics (Port)

Power over Ethernet

The NXA-ENET8-POE+ switch can provide DC power to a wide range of connected devices, eliminating the need for an additional power source and cutting down on the amount of cables attached to each device. Once configured to supply power, an automatic detection process is initialized by the switch that is authenticated by a PoE signature from the connected device.

Detection and authentication prevent damage to non-compliant devices (prior to IEEE 802.3af).

The switch's power management enables individual port power to be controlled within the switch's power budget. Port power can be automatically turned on and off for connected devices, and a per-port power priority can be set so that the switch never exceeds its power budget. When a device is connected to a switch port, its power requirements are detected by the switch before power is supplied. If the power required by a device exceeds the power budget of the port or the whole switch, power is not supplied.

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

NOTE: For more information on using the PoE provided by this switch refer to the *Installation* section on page 23.

Setting the Switch's Overall PoE Power Budget

Use the Administration > PoE > PSE (Configure Global) page to set the maximum PoE power budget for the switch (power available to all Gigabit Ethernet ports).

The following table lists the options on this page:

Administration - PoE (PSE) Options	
PoE Maximum Available Power	The power budget for the switch (i.e., power available to all switch ports). 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.
PoE Maximum Allocation Power	Sets a power budget for the switch. (Range: 50000-740000 milliwatts; Default: 370000 milliwatts)
Compatible Mode	<p>Allows the switch to detect and provide power to powered devices that were designed prior to the IEEE 802.3af PoE standard. (Default: Disabled)</p> <p>The switch automatically detects attached PoE devices by periodically transmitting test voltages that over the Gigabit Ethernet copper-media ports. When an IEEE 802.3af or 802.3at compatible device is plugged into one of these ports, the powered device reflects the test voltage back to the switch, which may then turn on the power to this device. When the compatibility mode is enabled, this switch can detect IEEE 802.3af or 802.3at compliant devices and the more recent 802.3af non-compliant devices that also reflect the test voltages back to the switch. It cannot detect other legacy devices that do not reflect back the test voltages.</p> <p>For legacy devices to be supported by this switch, they must be able to accept power over the data pairs connected to the RJ-45 ports.</p>

Perform these steps to set the overall PoE power budget for switch:

1. Click **Administration > PoE > PSE**.
2. Select **Configure Global** from the Step list.
3. Set the maximum PoE power provided by the switch, and enable the compatible mode if required.
4. Click **Apply**.

The screenshot shows the configuration page for PoE settings. At the top, it says 'Administration > PoE > PSE'. Below that, there is a 'Step:' dropdown menu set to '1. Configure Global'. The main configuration area has three rows:

- 'PoE Maximum Available Power' is set to '740000 milliwatts using joint power'.
- 'PoE Maximum Allocation Power (50000-740000)' has a text input field containing '370000' followed by 'milliwatts'.
- 'Compatible Mode' has a checked checkbox labeled 'Enabled'.

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

FIG. 233 Setting the Switch's PoE Budget

Setting the Port PoE Power Budget

Use the Administration > PoE > PSE 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. For the NXA-ENET8-POE+, the total PoE power delivered by all ports cannot exceed the maximum power budget of 160W.
The number of ports which can supply maximum power simultaneously to connected devices is listed in the following bullet points:
 - 30W (802.3at): 4
 - 15.4W (802.3af): 8
 - 7.5W (802.3af): 8
- If a device is connected to a switch port and the switch detects that it requires more than the power budget set for the port or to the overall switch, no power is supplied to the device (i.e., port power remains off).
- If the power demand from devices connected to all switch ports exceeds the power budget set for the switch, the port power priority settings are used to control the supplied power. For example:
 - If a device is connected to a low-priority port and causes the switch to exceed its budget, power to this port is not turned on.
 - If a device is connected to a critical or high-priority port and would cause the switch to exceed its power budget as determined during bootup, power is provided to the port only if the switch can drop power to one or more lower-priority ports and thereby remain within its overall budget.
 - If a device is connected to a port after the switch has finished booting up and would cause the switch to exceed its budget, power will not be provided to that port regardless of its priority setting.
 - If priority is not set for any ports, and 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 priority is not set for any ports, and PoE consumption exceeds the maximum power provided by the switch, power is shut down in the reverse sequence, starting from Port 24.

The following table lists the options on this page:

Administration - PoE (PSE) Options	
Port	The port number on the switch. (Range: 1-8/22/24/48)
Admin Status	Enables PoE power on a port. Power is automatically supplied when a device is detected on a port, providing that the power demanded does not exceed the switch or port power budget. (Default: Enabled)
Mode	Shows whether or not PoE power is being supplied to a port.
Time Range Time	Name of a time range. If a time range is set, then PoE will be provided to an interface during the specified period.
Time Range Status	Indicates if a time range has been applied to an interface, and whether it is currently active or inactive.
Priority	Sets the power priority for a port. (Options: Low, High or Critical; Default: Low)
Power Allocation	Sets the power budget for a port. (Range: 3000-30000 milliwatts; Default: 30000 milliwatts)
Power Consumption	Current power consumption on a port

Perform these steps to set the PoE power budget for a port:

1. Click **Administration > PoE > PSE**.
2. Enable PoE power on selected ports. Set the priority and the power budget. And specify a time range during which PoE will be provided to an interface.
3. Click **Apply**.

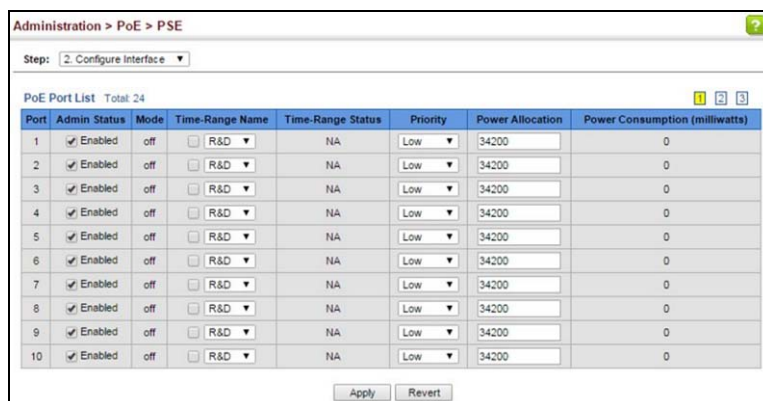


FIG. 234 Setting a Port's PoE Budget

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 on-board 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 on-board 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.

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	userdefined	userdefined	userdefined	userdefined	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	userdefined	userdefined	userdefined	userdefined	Community string only
v3	noAuthNoPriv	userdefined	userdefined	userdefined	userdefined	A username match only
v3	AuthNoPriv	userdefined	userdefined	userdefined	userdefined	Provides user authentication via MD5 or SHA algorithms
v3	AuthPriv	userdefined	userdefined	userdefined	userdefined	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.

The following table lists the options on this page:

Administration - SNMP Options	
Agent Status	Enables SNMP on the switch. (Default: Enabled)
Authentication Traps	Issues a notification message to specified IP trap managers whenever an invalid community string is submitted during the SNMP access authentication process. (Default: Enabled) 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.

Perform these steps 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**.

The screenshot shows the 'Administration > SNMP' configuration page. At the top, there is a breadcrumb 'Administration > SNMP' and a 'Step:' dropdown menu set to '1. Configure Global'. Below this, there are two rows of configuration options: 'Agent Status' and 'Authentication Traps'. Each row has a checkbox that is checked and the word 'Enabled' next to it. At the bottom right of the form, there are two buttons: 'Apply' and 'Revert'.

FIG. 235 Configuring Global Settings for SNMP

Setting the Local Engine ID

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

Command Usage

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

The following table lists the options on this page:

Administration - SNMP Options	
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 <i>123456789</i> is equivalent to <i>1234567890</i> .
Engine Boots	The number of times that the engine has (re-)initialized since the SNMP Engine ID was last configured.

Perform these steps 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**.

The screenshot shows the 'Administration > SNMP' configuration page. At the top, it indicates 'Step: 2. Configure Engine' and 'Action: Set Engine ID'. Below this, there are two input fields: 'Engine ID' with the value '800001030300000c0000fd0000' and 'Engine Boots' with the value '5'. At the bottom of the form, there are two buttons labeled 'Default' and 'Save'.

FIG. 236 Configuring the Local Engine ID for SNMP

Specifying a Remote Engine ID

Use the Administration > SNMP (Configure Engine - Add Remote Engine) page to configure an 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 the *Configuring Local SNMPv3 Users* section on page 215 for more information.)

The following table lists the options on this page:

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

Perform these steps 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 at least 9 hexadecimal characters, and the IP address of the remote host.
5. Click **Apply**.

FIG. 237 Configuring a Remote Engine ID for SNMP

Perform these steps 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.

FIG. 238 Showing Remote Engine IDs for SNMP

Setting SNMPv3 Views

Use the Administration > SNMP (Configure View) page to configure SNMPv3 views which are used to restrict user access to specified portions of the MIB tree. The predefined view *defaultview* includes access to the entire MIB tree.

The following table lists the options on this page:

Administration - SNMP Options	
Add View	
View Name	The name of the SNMP view. (Range: 1-32 characters)
OID Subtree	Specifies the initial object identifier of a branch within the MIB tree. Wild cards can be used to mask a specific portion of the OID string. Use the Add OID Subtree page to configure additional object identifiers. (Range: 1-64 characters)
Type	Indicates if the object identifier of a branch within the MIB tree is included or excluded from the SNMP view.
Add OID Subtree	
View Name	Lists the SNMP views configured in the Add View page. (Range: 1-32 characters)
OID Subtree	Adds an additional object identifier of a branch within the MIB tree to the selected View. Wild cards can be used to mask a specific portion of the OID string. (Range: 1-64 characters)
Type	Indicates if the object identifier of a branch within the MIB tree is included or excluded from the SNMP view.

Perform these steps 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**.

The screenshot shows the 'Administration > SNMP' page. At the top, there are two dropdown menus: 'Step: 3. Configure View' and 'Action: Add View'. Below these are three input fields: 'View Name' with the value 'ifEntry.a', 'OID Subtree' with the value '1.3.6.1.2.1.2.2.1.1.*', and 'Type' with the value 'Included'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

FIG. 239 Creating an SNMP View

Perform these steps 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.

The screenshot shows the 'Administration > SNMP' page. At the top, there are two dropdown menus: 'Step: 3. Configure View' and 'Action: Show View'. Below these is a table titled 'SNMPv3 View List' with a 'Total: 2' indicator. The table has two columns: a checkbox and 'View Name'. There are two rows: one with a checked checkbox and 'ifEntry.a', and another with an unchecked checkbox and 'defaultview'. At the bottom right, there are two buttons: 'Delete' and 'Revert'.

View Name
ifEntry.a
defaultview

FIG. 240 Showing SNMP Views

Perform these steps 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**.

The screenshot shows the 'Administration > SNMP' page. At the top, there are two dropdown menus: 'Step: 3. Configure View' and 'Action: Add OID Subtree'. Below these are three input fields: 'View Name' with a dropdown menu showing 'ifEntry.a', 'OID Subtree' with the value '1.3.6.1.2.1.2.2.1.2.*', and 'Type' with the value 'Included'. At the bottom right, there are two buttons: 'Apply' and 'Revert'.

FIG. 241 Adding an OID Subtree to an SNMP View

Perform these steps 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.

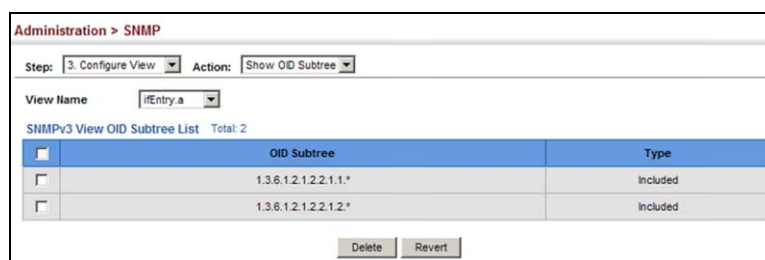


FIG. 242 Showing the OID Subtree Configured for SNMP Views

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.

The following table lists the options on this page:

Administration - SNMP Options	
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: <ul style="list-style-type: none"> • noAuthNoPriv - There is no authentication or encryption used in SNMP communications. (This is the default security level.) • AuthNoPriv - SNMP communications use authentication, but the data is not encrypted. • AuthPriv - SNMP communications use both authentication and encryption.
Read View	The configured view for read access. (Range: 1-32 characters)
Write View	The configured view for write access. (Range: 1-32 characters)
Notify View	The configured view for notifications. (Range: 1-32 characters)

Supported Notification Messages		
Model	Level	Group
<i>RFC 1493 Traps</i>		
newRoot	1.3.6.1.2.1.17.0.1	The newRoot trap indicates that the sending agent has become the new root of the Spanning Tree; the trap is sent by a bridge soon after its election as the new root, e.g., upon expiration of the Topology Change Timer immediately subsequent to its election.
topologyChange	1.3.6.1.2.1.17.0.2	A topologyChange trap is sent by a bridge when any of its configured ports transitions from the Learning state to the Forwarding state, or from the Forwarding state to the Discarding state. The trap is not sent if a newRoot trap is sent for the same transition.
<i>SNMPv2 Traps</i>		
coldStart	1.3.6.1.6.3.1.1.5.1	A coldStart trap signifies that the SNMPv2 entity, acting in an agent role, is reinitializing itself and that its configuration may have been altered.
warmStart	1.3.6.1.6.3.1.1.5.2	A warmStart trap signifies that the SNMPv2 entity, acting in an agent role, is reinitializing itself such that its configuration is unaltered.
linkDown*	1.3.6.1.6.3.1.1.5.3	A linkDown trap signifies that the SNMP entity, acting in an agent role, has detected that the ifOperStatus object for one of its communication links is about to enter the down state from some other state (but not from the notPresent state). This other state is indicated by the included value of ifOperStatus.

Supported Notification Messages		
linkUp*	1.3.6.1.6.3.1.1.5.4	A linkUp trap signifies that the SNMP entity, acting in an agent role, has detected that the ifOperStatus object for one of its communication links left the down state and transitioned into some other state (but not into the notPresent state). This other state is indicated by the included value of ifOperStatus.
authenticationFailure*	1.3.6.1.6.3.1.1.5.5	An authenticationFailure trap signifies that the SNMPv2 entity, acting in an agent role, has received a protocol message that is not properly authenticated. While all implementations of the SNMPv2 must be capable of generating this trap, the snmpEnableAuthenTraps object indicates whether this trap will be generated.
<i>RMONEvents(V2)</i>		
risingAlarm	1.3.6.1.2.1.16.0.1	The SNMP trap that is generated when an alarm entry crosses its rising threshold and generates an event that is configured for sending SNMP traps.
fallingAlarm	1.3.6.1.2.1.16.0.2	The SNMP trap that is generated when an alarm entry crosses its falling threshold and generates an event that is configured for sending SNMP traps.
<i>Private Traps</i>		
swPowerStatusChangeTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.1	This trap is sent when the power state changes.
swPortSecurityTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.36	This trap is sent when the port is being intruded. This trap will only be sent when the portSecActionTrap is enabled.
swIpFilterRejectTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.40	This trap is sent when an incorrect IP address is rejected by the IP Filter.
pethPsePortOnOffNotification	1.3.6.1.4.1.259.6.10.120.2.1.0.43	This notification indicates if a PSE port is delivering power to the PD. This notification should be sent on every status change except in searching mode.
pethPsePortPowerMaintenanceStatus Notification	1.3.6.1.4.1.259.6.10.120.2.1.0.44	This notification indicates a Port Change Status and should be sent on every status change.
pethMainPowerUsageOnNotification	1.3.6.1.4.1.259.6.10.120.2.1.0.45	This notification indicates PSE threshold usage indication is on; and the power usage is above the threshold.
pethMainPowerUsageOffNotification	1.3.6.1.4.1.259.6.10.120.2.1.0.46	This notification indicates that the PSE threshold usage indication is off; and the usage power is below the threshold.
swAtcBcastStormAlarmFireTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.70	When broadcast traffic is detected as a storm, this trap is fired.
swAtcBcastStormAlarmClearTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.71	When a broadcast storm is detected as normal traffic, this trap is fired.
swAtcBcastStormTcApplyTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.72	When ATC is activated, this trap is fired.
swAtcBcastStormTcReleaseTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.73	When ATC is released, this trap is fired.
swAtcMcastStormAlarmFireTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.74	When multicast traffic is detected as the storm, this trap is fired.
swAtcMcastStormAlarmClearTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.75	When multicast storm is detected as normal traffic, this trap is fired.
swAtcMcastStormTcApplyTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.76	When ATC is activated, this trap is fired.
swAtcMcastStormTcReleaseTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.77	When ATC is released, this trap is fired.
stpBpduGuardPortShutdownTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.91	This trap will be sent when an interface is shut down because of BPDU guard.
swLoopbackDetectionTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.95	This trap is sent when loopback BPDUs have been detected.
networkAccessPortLinkDetectionTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.96	This trap is sent when a networkAccessPortLinkDetection event is triggered.
dot1agCfmMepUpTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.97	This trap is sent when a new remote MEP is discovered.
dot1agCfmMepDownTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.98	This trap is sent when port status or interface status TLV received from remote MEP indicates it is not up.
dot1agCfmConfigFailTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.99	This trap is sent when a MEP receives a CCM with MPID which already exists on the same MA in this switch.
dot1agCfmLoopFindTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.100	This trap is sent when a MEP receives its own CCMs.

Supported Notification Messages		
dot1agCfmMepUnknownTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.101	This trap is sent when a CCM is received from an unexpected MEP.
dot1agCfmMepMissingTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.102	This trap is sent when the cross-check enable timer expires and no CCMs were received from an expected (configured) MEP.
dot1agCfmMaUpTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.103	This trap is sent when all expected remote MEPs are up.
autoUpgradeTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.104	This trap is sent when auto upgrade is executed.
swCpuUtiRisingNotification	1.3.6.1.4.1.259.6.10.120.2.1.0.107	This notification indicates that the CPU utilization has risen from cpuUtiFallingThreshold to cpuUtiRisingThreshold.
swCpuUtiFallingNotification	1.3.6.1.4.1.259.6.10.120.2.1.0.108	This notification indicates that the CPU utilization has fallen from cpuUtiRisingThreshold to cpuUtiFallingThreshold.
swMemoryUtiRisingThreshold Notification	1.3.6.1.4.1.259.6.10.120.2.1.0.109	This notification indicates that the memory utilization has risen from memoryUtiFallingThreshold to memoryUtiRisingThreshold.
swMemoryUtiFallingThreshold Notification	1.3.6.1.4.1.259.6.10.120.2.1.0.110	This notification indicates that the memory utilization has fallen from memoryUtiRisingThreshold to memoryUtiFallingThreshold.
dhcpRougeServerAttackTrap	1.3.6.1.4.1.259.10.1.24.2.1.0.114	This trap is sent when receiving a DHCP packet from a rouge server.
macNotificationTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.138	This trap is sent when there are changes of the dynamic MAC addresses on the switch.
lbdDetectionTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.141	This trap is sent when a loopback condition is detected by LBD.
lbdRecoveryTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.142	This trap is sent when a recovery is done by LBD.
sfpThresholdAlarmWarnTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.189	This trap is sent when the SFP's A/D quantity is not within alarm/warning thresholds.
udldPortShutdownTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.192	This trap is sent when the port is shut down by UDLD.
userAuthenticationFailureTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.199	This trap will be triggered if authentication fails.
userAuthenticationSuccessTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.200	This trap will be triggered if authentication is successful.
loginTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.201	This trap is sent when user logs in.
logoutTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.202	This trap is sent when user logs out.
fileCopyTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.208	This trap is sent when file copy is executed. If the copy action is triggered by the system, the login user information (trapVarLoginUserName/ trapVarSessionType/ trapVarLoginInetAddressTypes/ trapVarLoginInetAddress) will be a null value.
userauthCreateUserTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.209	This trap is sent when a user account is created.
userauthDeleteUserTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.210	This trap is sent when a user account is deleted.
userauthModifyUserPrivilegeTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.211	This trap is sent when user privilege is modified.
cpuGuardControlTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.213	This trap is sent when CPU utilization rises above the high-watermark the first time or when CPU utilization rises from below the low-watermark to above the high-watermark.
cpuGuardReleaseTrap	1.3.6.1.4.1.259.6.10.120.2.1.0.214	This trap is sent when CPU utilization falls from above the high-watermark to below the low-watermark.

* These are legacy notifications and therefore must be enabled in conjunction with the corresponding traps on the SNMP Configuration menu.

Perform these steps 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**.

FIG. 243 Creating an SNMP Group

Perform these steps to show SNMP groups:

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

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

FIG. 244 Showing SNMP Groups

Setting Community Access Strings

Use the Administration > SNMP (Configure Community - Add) 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. The following table lists the options on this page:

Administration - SNMP Options	
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: <ul style="list-style-type: none"> • 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.

Perform these steps to set a community access string:

1. Click **Administration > SNMP**.
2. Select **Configure Community** from the Step list.
3. Select **Add** from the Action list.
4. Add new community strings as required, and select the corresponding access rights from the Access Mode list.
5. Click **Apply**.

Administration > SNMP

Step: 2. Configure Community Action: Add

Community String: spiderman

Access Mode: Read/Write

Apply Revert

FIG. 245 Setting Community Access Strings

Perform these steps to show the community access strings:

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

Administration > SNMP

Step: 2. Configure Community Action: Show

SNMP Community String List Total: 3

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

Delete Revert

FIG. 246 Showing Community Access Strings

Configuring Local SNMPv3 Users

Use the Administration > SNMP (Configure User - Add SNMPv3 Local User) page to authorize management access for SNMPv3 clients, or to identify the source of SNMPv3 trap messages sent from the local switch. Each SNMPv3 user is defined by a unique name. Users must be configured with a specific security level and assigned to a group. The SNMPv3 group restricts users to a specific read, write, and notify view.

The following table lists the options on this page:

Administration - SNMP Options	
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: <ul style="list-style-type: none"> • noAuthNoPriv - There is no authentication or encryption used in SNMP communications. (This is the default security level.) • AuthNoPriv - SNMP communications use authentication, but the data is not encrypted. • AuthPriv - SNMP communications use both authentication and encryption.
Authentication Protocol	The method used for user authentication. (Options: MD5, SHA; Default: MD5)
Authentication Password	A minimum of eight plain text characters is required. (Range: 8-32 characters)
Privacy Protocol	The encryption algorithm use for data privacy; only 56-bit DES is currently available.
Privacy Password	A minimum of eight plain text characters is required.

Perform these steps 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**.

Administration > SNMP

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

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:

FIG. 247 Configuring Local SNMPv3 Users

Perform these steps 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.

Administration > SNMP

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

SNMPv3 Local User List Total: 1

	User Name	Group Name	Model	Level	Authentication	Privacy
<input type="checkbox"/>	chris	r&d	v3	authPriv	MD5	DES56

FIG. 248 Showing Local SNMPv3 Users

Perform these steps to change a local SNMPv3 local user group:

1. Click **Administration > SNMP**.
2. Select **Change SNMPv3 Local User Group** from the Action list.
3. Select the User Name.
4. Enter a new group name.
5. Click **Apply**.

Administration > SNMP

Step: 5. Configure User Action: Change SNMPv3 Local User Group

User Name:

Group Name: bart public

FIG. 249 Changing a Local SNMPv3 User Group

Configuring Remote SNMPv3 Users

Use the Administration > SNMP (Configure User - Add SNMPv3 Remote User) page to identify the source of SNMPv3 inform messages sent from the local switch. Each SNMPv3 user is defined by a unique name. Users must be configured with a specific security level and assigned to a group. The SNMPv3 group restricts users to a specific read, write, and notify view.

Command Usage

To grant management access to an SNMPv3 user on a remote device, you must first specify the engine identifier for the SNMP agent on the remote device where the user resides. The remote engine ID is used to compute the security digest for authentication and encryption of packets passed between the switch and the remote user. (See the *Specifying Trap Managers* section on page 218 and the *Specifying a Remote Engine ID* section on page 208.)

The following table lists the options on this page:

Administration - SNMP Options	
User Name	The name of user connecting to the SNMP agent. (Range: 1-32 characters)
Group Name	The name of the SNMP group to which the user is assigned. (Range: 1-32 characters)
Remote IP	IPv4 address of the remote device where the user resides.
Security Model	The user security model; SNMP v1, v2c or v3.
Security Level	The following security levels are only used for the groups assigned to the SNMP security model: <ul style="list-style-type: none"> 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.

Perform these steps 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 protocol and password must also be specified.
5. Click **Apply**.

The screenshot shows the configuration page for a remote SNMPv3 user. The breadcrumb is 'Administration > SNMP'. The 'Step' is '5. Configure User' and the 'Action' is 'Add SNMPv3 Remote User'. The form is titled 'SNMPv3 User' and contains the following fields:

- User Name:** mark
- Group Name:** public (selected) and r&d
- Remote IP:** 192.168.1.19
- Security Model:** v3
- Security Level:** authPriv
- User Authentication:**
 - Authentication Protocol:** MD5
 - Authentication Password:** greenpeace
- Data Privacy:**
 - Privacy Protocol:** DES56
 - Privacy Password:** einstien

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

FIG. 250 Configuring Remote SNMPv3 Users

Perform these step 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.

Administration > SNMP							
Step: 5. Configure User		Action: Show SNMPv3 Remote User					
SNMPv3 Remote User List Total: 1							
	User Name	Group Name	Engine ID	Model	Level	Authentication	Privacy
<input type="checkbox"/>	mark	r&d	5432100000	v3	authPriv	MD5	DES56
<input type="button" value="Delete"/> <input type="button" value="Revert"/>							

FIG. 251 Showing Remote SNMPv3 Users

Specifying Trap Managers

Use the Administration > SNMP (Configure Trap) page to specify the host devices to be sent traps and the types of traps to send. Traps indicating status changes are issued by the switch to the specified trap managers. You must specify trap managers so that key events are reported by this switch to your management station (using network management software). You can specify up to five management stations that will receive authentication failure messages and other trap messages from the switch.

Command Usage

Notifications are issued by the switch as trap messages by default. The recipient of a trap message does not send a response to the switch. Traps are therefore not as reliable as inform messages, which include a request for acknowledgment of receipt. Informs can be used to ensure that critical information is received by the host. However, note that informs consume more system resources because they must be kept in memory until a response is received. Informs also add to network traffic. You should consider these effects when deciding whether to issue notifications as traps or informs.

To send an inform to a SNMPv2c host, complete these steps:

1. Enable the SNMP agent (page 207).
2. Create a view with the required notification messages (page 209).
3. Configure the group (matching the community string specified on the Configure Trap - Add page) to include the required notify view (page 211).
4. Enable trap informs as described in the following pages.

To send an inform to a SNMPv3 host, complete these steps:

1. Enable the SNMP agent (page 208).
2. Create a remote SNMPv3 user to use in the message exchange process (page 215). If the user specified in the trap configuration page does not exist, an SNMPv3 group will be automatically created using the name of the specified remote user, and default settings for the read, write, and notify view.
3. Create a view with the required notification messages (page 209).
4. Create a group that includes the required notify view (page 211).
5. Enable trap informs as described in the following pages.

The following table lists the options on this page:

Administration - SNMP Options	
SNMP Version 1	
IP Address	IPv4 or IPv6 address of a new management station to receive notification message (i.e., the targeted recipient).
Version	Specifies whether to send notifications as SNMP v1, v2c, or v3 traps. (Default: v1)
Community String	Specifies a valid community string for the new trap manager entry. (Range: 1-32 characters, case sensitive) Although you can set this string in the Configure Trap - Add page, we recommend defining it in the Configure User - Add Community page.
UDP Port	Specifies the UDP port number used by the trap manager. (Default: 162)
SNMP Version 2c	
IP Address	IPv4 or IPv6 address of a new management station to receive notification message (i.e., the targeted recipient).
Version	Specifies whether to send notifications as SNMP v1, v2c, or v3 traps.

Administration - SNMP Options	
Notification Type	<ul style="list-style-type: none"> Traps - Notifications are sent as trap messages. Inform - Notifications are sent as inform messages. Note that this option is only available for version 2c and 3 hosts. (Default: traps are used) Timeout - The number of seconds to wait for an acknowledgment before resending an inform message. (Range: 0-2147483647 centiseconds; Default: 1500 centiseconds) Retry times - The maximum number of times to resend an inform message if the recipient does not acknowledge receipt. (Range: 0-255; Default: 3)
Community String	Specifies a valid community string for the new trap manager entry. (Range: 1-32 characters, case sensitive) Although you can set this string in the Configure Trap - Add page, we recommend defining it in the Configure User - Add Community page.
UDP Port	Specifies the UDP port number used by the trap manager. (Default: 162)
SNMP Version 3	
IP Address	IPv4 or IPv6 address of a new management station to receive notification message (i.e., the targeted recipient).
Version	Specifies whether to send notifications as SNMP v1, v2c, or v3 traps.
Notification Type	<ul style="list-style-type: none"> Traps - Notifications are sent as trap messages. Inform - Notifications are sent as inform messages. Note that this option is only available for version 2c and 3 hosts. (Default: traps are used) Timeout - The number of seconds to wait for an acknowledgment before resending an inform message. (Range: 0-2147483647 centiseconds; Default: 1500 centiseconds) Retry times - The maximum number of times to resend an inform message if the recipient does not acknowledge receipt. (Range: 0-255; Default: 3)
Local User Name	The name of a local user which is used to identify the source of SNMPv3 trap messages sent from the local switch. (Range: 1-32 characters) If an account for the specified user has not been created (page 363), 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 365), one will be automatically generated.
UDP Port	Specifies the UDP port number used by the trap manager. (Default: 162)
Security Level	When trap version 3 is selected, you must specify one of the following security levels. (Default: noAuthNoPriv) <ul style="list-style-type: none"> 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.

Perform these steps to configure trap managers:

1. Click **Administration > SNMP**.
2. Select **Configure Trap** from the Step list.
3. Select **Add** from the Action list.
4. Fill in the required parameters based on the selected SNMP version.
5. Click **Apply**.

The screenshot shows the 'Administration > SNMP' configuration interface. At the top, it indicates the current step is '6. Configure Trap' and the action is 'Add'. Below this, there are four input fields: 'IP Address' with the value '192.168.0.3', 'Version' with a dropdown menu set to 'v1', 'Community String' with the value 'private', and 'UDP Port (1-65535)' with the value '162'. At the bottom right of the form, there are two buttons: 'Apply' and 'Revert'.

FIG. 252 Configuring Trap Managers (SNMPv1)

Administration > SNMP

Step: 6. Configure Trap Action: Add

IP Address: 192.168.2.9

Version: v2c

Notification Type: Inform

Timeout (0-2147483647): centiseconds

Retry Times (0-255):

Community String: venus

UDP Port (1-65535):

Apply Revert

FIG. 253 Configuring Trap Managers (SNMPv2c)

Administration > SNMP

Step: 6. Configure Trap Action: Add

IP Address: 192.168.2.9

Version: v3

Notification Type: Inform

Timeout (0-2147483647): centiseconds

Retry Times (0-255):

Remote User Name:

UDP Port (1-65535):

Security Level: authPriv

Apply Revert

FIG. 254 Configuring Trap Managers (SNMPv3)

Perform these steps to show configured trap managers:

1. Click **Administration > SNMP**.
2. Select **Configure Trap** from the Step list.
3. Select **Show** from the Action list.

Administration > SNMP

Step: 6. Configure Trap Action: Show

SNMP Trap Manager List Total: 5

<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

Delete Revert

FIG. 255 Showing Trap Managers

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 43. 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 218, a default notify filter will be created.

The following table lists the options on this page:

Administration - SNMP Options	
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)

Perform these steps 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**.

FIG. 256 Creating SNMP Notification Logs

Perform these steps 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.

FIG. 257 Showing SNMP Notification Logs

Showing SNMP Statistics

Use the Administration > SNMP (Show Statistics) page to show counters for SNMP input and output protocol data units. The following table lists the options on this page:

Administration - SNMP Options	
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 <i>tooBig</i> .
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 <i>noSuchName</i> .
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 <i>badValue</i> .
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 <i>genErr</i> .
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.

Perform these steps to show SNMP statistics:

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

Administration > SNMP			
Step: 8. Show Statistics			
SNMP Statistics			
SNMP packets Input	0	SNMP packets Output	0
Bad SNMP version errors	0	Too big errors	0
Unknown community name	0	No such name errors	0
Illegal operation for community name supplied	0	Bad values errors	0
Encoding errors	0	General errors	0
Number of requested variables	0	Response PDUs	0
Number of altered variables	0	Trap PDUs	0
Get-request PDUs	0		
Get-next PDUs	0		
Set-request PDUs	0		

Refresh

FIG. 258 Showing SNMP Statistics

Remote Monitoring

Remote Monitoring allows a remote device to collect information or respond to specified events on an independent basis. This switch is an RMON-capable device which can independently perform a wide range of tasks, significantly reducing network management traffic. It can continuously run diagnostics and log information on network performance. If an event is triggered, it can automatically notify the network administrator of a failure and provide historical information about the event. If it cannot connect to the management agent, it will continue to perform any specified tasks and pass data back to the management station the next time it is contacted.

The switch supports mini-RMON, which consists of the Statistics, History, Event and Alarm groups. When RMON is enabled, the system gradually builds up information about its physical interfaces, storing this information in the relevant RMON database group. A management agent then periodically communicates with the switch using the SNMP protocol. However, if the switch encounters a critical event, it can automatically send a trap message to the management agent which can then respond to the event if so configured.

Configuring RMON Alarms

Use the Administration > RMON (Configure Global - Add - Alarm) page to define specific criteria that will generate response events. Alarms can be set to test data over any specified time interval, and can monitor absolute or changing values (such as a statistical counter reaching a specific value, or a statistic changing by a certain amount over the set interval). Alarms can be set to respond to rising or falling thresholds. (However, note that after an alarm is triggered it will not be triggered again until the statistical value crosses the opposite bounding threshold and then back across the trigger threshold.

Command Usage

If an alarm is already defined for an index, the entry must be deleted before any changes can be made.

The following table lists the options on this page:

Administration - RMON Options	
Index	Index to this entry. (Range: 1-65535)
Variable	The object identifier of the MIB variable to be sampled. Only variables of the type etherStatsEntry.n.n may be sampled. Note that etherStatsEntry.n uniquely defines the MIB variable, and etherStatsEntry.n.n defines the MIB variable, plus the etherStatsIndex. For example, 1.3.6.1.2.1.16.1.1.1.6.1 denotes etherStatsBroadcastPkts, plus the etherStatsIndex of 1.
Interval	The polling interval. (Range: 1-31622400 seconds)
Sample Type	Tests for absolute or relative changes in the specified variable. <ul style="list-style-type: none"> Absolute - The variable is compared directly to the thresholds at the end of the sampling period. Delta - The last sample is subtracted from the current value and the difference is then compared to the thresholds.
Rising Threshold	If the current value is greater than or equal to the rising threshold, and the last sample value was less than this threshold, then an alarm will be generated. After a rising event has been generated, another such event will not be generated until the sampled value has fallen below the rising threshold, reaches the falling threshold, and again moves back up to the rising threshold. (Range: 0-2147483647)
Rising Event Index	The index of the event to use if an alarm is triggered by monitored variables reaching or crossing above the rising threshold. If there is no corresponding entry in the event control table, then no event will be generated. (Range: 0-65535)
Falling Threshold	If the current value is less than or equal to the falling threshold, and the last sample value was greater than this threshold, then an alarm will be generated. After a falling event has been generated, another such event will not be generated until the sampled value has risen above the falling threshold, reaches the rising threshold, and again moves back down to the falling threshold. (Range: 0-2147483647)
Falling Event Index	The index of the event to use if an alarm is triggered by monitored variables reaching or crossing below the falling threshold. If there is no corresponding entry in the event control table, then no event will be generated. (Range: 0-65535)
Owner	Name of the person who created this entry. (Range: 1-32 characters)

Perform these steps 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**.

FIG. 259 Configuring an RMON Alarm

Perform these steps 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**.

Index	Status	Variable	Interval	Type	Last Value	Rising Threshold	Rising Event Index	Falling Threshold	Falling Event Index	Owner
1	Valid	1.3.6.1.2.1.16.1.1.1.6.1	30	Delta	0	892800	0	446400	0	
2	Valid	1.3.6.1.2.1.16.1.1.1.6.2	30	Delta	0	892800	0	446400	0	
3	Valid	1.3.6.1.2.1.16.1.1.1.6.3	30	Delta	0	892800	0	446400	0	
4	Valid	1.3.6.1.2.1.16.1.1.1.6.4	30	Delta	0	892800	0	446400	0	
5	Valid	1.3.6.1.2.1.16.1.1.1.6.5	30	Delta	0	892800	0	446400	0	

FIG. 260 Showing Configured RMON Alarms

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

The following table lists the options on this page:

Administration - RMON Options	
Index	Index to this entry. (Range: 1-65535)
Type	Specifies the type of event to initiate: <ul style="list-style-type: none"> • 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 the <i>System Log Configuration</i> section on page 187). • Trap - Sends a trap message to all configured trap managers (see the <i>Specifying Trap Managers</i> section on page 218). • 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 the <i>Setting Community Access Strings</i> section on page 214) prior to configuring it here. (Range: 1-32 characters)
Description	A comment that describes this event. (Range: 1-127 characters)
Owner	Name of the person who created this entry. (Range: 1-32 characters)

Perform these steps 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**.

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' and 'Event', with 'Event' selected. The form contains several input fields: 'Index (1-65535)' with the value '2', 'Type' with a dropdown menu set to 'Log and Trap', 'Community' with the text 'private', 'Description' with the text 'for software group', and 'Owner' with the text 'david'. At the bottom right of the form are two buttons: 'Apply' and 'Revert'.

FIG. 261 Configuring an RMON Event

Perform these steps 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**.

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

FIG. 262 Showing Configured RMON Events

Configuring RMON History Samples

Use the Administration > RMON (Configure Interface - Add - History) page to collect statistics on a physical interface to monitor network utilization, packet types, and errors. A historical record of activity can be used to track down intermittent problems. The record can be used to establish normal baseline activity, which may reveal problems associated with high traffic levels, broadcast storms, or other unusual events. It can also be used to predict network growth and plan for expansion before your network becomes too overloaded.

Command Usage

- Each index number equates to a port on the switch.
- If history collection is already enabled on an interface, the entry must be deleted before any changes can be made.
- The information collected for each sample includes:
input octets, packets, broadcast packets, multicast packets, undersized 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 the *Showing Port or Trunk Statistics* section on page 63.
- 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.

The following table lists the options on this page:

Administration - RMON Options	
Port	The port number on the switch. (Range: 1-10/26/28/52)
Index	Index to this entry. (Range: 1-65535)
Interval	The polling interval. (Range: 1-3600 seconds; Default: 1800 seconds)
Buckets	The number of buckets requested for this entry. (Range: 1-65536; Default: 8) The number of buckets granted are displayed on the Show page.
Owner	Name of the person who created this entry. (Range: 1-32 characters)

Perform these steps 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**.

FIG. 263 Configuring an RMON History Sample

Perform these steps 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**.

<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	

FIG. 264 Showing Configured RMON History Samples

Perform these steps 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**.

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

FIG. 265 Showing Collected RMON History Samples

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.

The following table lists the options on this page:

Administration - RMON Options	
Port	The port number on the switch. (Range: 1-10/26/28/52)
Index	Index to this entry. (Range: 1-65535)
Owner	Name of the person who created this entry. (Range: 1-32 characters)

Perform these steps 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**.

Administration > RMON

Step: 2. Configure Interface Action: Add

History Statistics

Port 2

Index (1-65535) 100

Owner mary

Apply Revert

FIG. 266 Configuring an RMON Statistical Sample

Perform these steps 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**.

Administration > RMON

Step: 2. Configure Interface Action: Show

History Statistics

Port 2

RMON Statistics Port List Total: 2

	Index	Status	Owner
<input type="checkbox"/>	1	Valid	abc
<input type="checkbox"/>	2	Valid	test

Delete Revert

FIG. 267 Showing Configured RMON Statistical Samples

Perform these steps 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**.

RMON Statistics Port Details			
Received Octets	9613105	Collisions	0
Received Packets	24621	Drop Events	0
Broadcast Packets	608	Frames of 64 Octets	13595
Multicast Packets	5538	Frames of 65 to 127 Octets	2606
Undersize Packets	0	Frames of 128 to 255 Octets	1222
Oversize Packets	0	Frames of 256 to 511 Octets	56
CRC Align Errors	0	Frames of 512 to 1023 Octets	2028
Jabbers	0	Frames of 1024 to 1518 Octets	5114
Fragments	0		

FIG. 268 Showing Collected RMON Statistical Samples

Setting a Time Range

Use the Administration > Time Range page to set a time range during which various functions are applied, including applied ACLs or PoE.

Command Usage

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

The following table lists the options on this page:

Administration - Time Range Options	
Add	
Time-Range Name	Name of a time range. (Range: 1-32 characters)
Add Rule	
Time-Range	Name of a time range.
Mode	<ul style="list-style-type: none"> • 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.

Perform these steps to configure a time range:

1. Click **Administration > Time Range**.
2. Select **Add** from the Action list.
3. Enter the name of a time range.
4. Click **Apply**.

FIG. 269 Setting the Name of a Time Range

Perform these steps to show a list of time ranges:

1. Click **Administration > Time Range**.
2. Select **Show** from the Action list.



FIG. 270 Showing a List of Time Ranges

Perform these steps to configure a rule for a time range:

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



FIG. 271 Add a Rule to a Time Range

Perform these steps to show the rules configured for a time range:

1. Click **Administration > Time Range**.
2. Select **Show Rule** from the Action list.




FIG. 272 Showing the Rules Configured for a Time Range

LBD Configuration

The switch can be configured to detect general loopback conditions caused by hardware problems or faulty protocol settings. When enabled, a control frame is transmitted on the participating ports, and the switch monitors inbound traffic to see if the frame is looped back.

Usage Guidelines

- The default settings for the control frame transmit interval and recover time may be adjusted to improve performance for your specific environment. The shutdown mode may also need to be changed once you determine what kind of packets are being looped back.
- General loopback detection provided by the commands described in this section and loopback detection provided by the spanning tree protocol cannot both be enabled at the same time. If loopback detection is enabled for the spanning tree protocol, general loopback detection cannot be enabled on the same interface.
- When a loopback event is detected on an interface or when an interface is released from a shutdown state caused by a loopback event, a trap message is sent and the event recorded in the system log.
- Loopback detection must be enabled both globally and on an interface for loopback detection to take effect.

Configuring Global Settings for LBD

Use the Administration > LBD (Configure Global) page to enable loopback detection globally, specify the interval at which to transmit control frames, the interval to wait before releasing an interface from shutdown state, the response to a detected loopback, and the traps to send.

The following table lists the options on this page:

Administration - LBD (Configure Global) Options	
Global Status	Enables loopback detection globally on the switch. (Default: Enabled)
Transmit Interval	Specifies the interval at which to transmit loopback detection control frames. (Range: 1-32767 seconds; Default: 10 seconds)
Recover Time	Specifies the interval to wait before the switch automatically releases an interface from shutdown state. (Range: 60-1,000,000 seconds; Default: 60 seconds) When the loopback detection mode is changed, any ports placed in shutdown state by the loopback detection process will be immediately restored to operation regardless of the remaining recover time. If the recover time is not enabled (checkbox unmarked), all ports placed in shutdown state can be restored to operation using the Release button. To restore a specific port, re-enable Admin status on the Configure Interface page.
Action	Specifies the protective action the switch takes when a loopback condition is detected. (Options: None, Shutdown; Default: Shutdown) <ul style="list-style-type: none"> • None - No action is taken. • Shutdown - When the response to a detected loopback condition is set to shut down a port, and a port receives a control frame sent by itself, this means that the port is in looped state, and the VLAN in the frame payload is also in looped state with the wrong VLAN tag. The looped port is therefore shut down. When the loopback detection response is changed, any ports placed in shutdown state by the loopback detection process will be immediately restored to operation regardless of the remaining recover time.
Trap	Sends a trap when a loopback condition is detected, or when the switch recovers from a loopback condition. (Options: Both, Detect, None, Recover; Default: None) <ul style="list-style-type: none"> • Both - Sends an SNMP trap message when a loopback condition is detected, or when the switch recovers from a loopback condition. • Detect - Sends an SNMP trap message when a loopback condition is detected. • None - Does not send an SNMP trap for loopback detection or recovery. • Recover - Sends an SNMP trap message when the switch recovers from a loopback condition.
Release	Releases all interfaces currently shut down by the loopback detection feature.

Perform these steps to configure global settings for LBD:

1. Click **Administration > LBD > Configure Global**.
2. Make the required configuration changes.
3. Click **Apply**.

Administration > LBD

Step: 1. Configure Global

Global Status Enabled

Transmit Interval (1-32767) sec

Recover Time (60-1000000) sec

Action

Trap

Click this button to release all looped ports manually

FIG. 273 Configuring Global Settings for LBD

Configuring Interface Settings for LBD

Use the Administration > LBD (Configure Interface) page to enable loopback detection on an interface, to display the loopback operational state, and the VLANs which are looped back.

The following table lists the options on this page:

Administration - LBD (Configure Interface) Options	
Interface	Displays a list of ports or trunks. <ul style="list-style-type: none"> • Port - Port Identifier (Range: 1-10/26/28/52) • Trunk - Trunk Identifier (Range: 1-8)
Admin State	Manually enables or disables an interface. (Default: Enabled)
Operation State	Valid states include Normal or Looped.
Looped VLAN	Shows the VLANs which are in looped state

Perform these steps to configure interface settings for LBD:

1. Click **Administration > LBD > Configure Interface**.
2. Make the required configuration changes.
3. Click **Apply**.

Administration > LBD

Step: 2. Configure Interface

Interface Port Trunk

Port List Total: 28

Port	Admin State	Operation State	Looped VLAN
1	<input checked="" type="checkbox"/> Enabled	Normal	None
2	<input checked="" type="checkbox"/> Enabled	Normal	None
3	<input checked="" type="checkbox"/> Enabled	Normal	None
4	<input checked="" type="checkbox"/> Enabled	Normal	None
5	<input checked="" type="checkbox"/> Enabled	Normal	None

FIG. 274 Configuring Interface Settings for LBD

Multicast Filtering

This chapter describes how to configure the following multicast services:

- **IGMP Snooping** - Configures snooping and query parameters.
- **Filtering and Throttling** - Filters specified multicast service, or throttles the maximum of multicast groups allowed on an interface.
- **MLD Snooping** - Configures snooping and query parameters for IPv6.

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.




FIG. 275 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.

Layer 2 IGMP (Snooping and Query for IPv4)

IGMP Snooping and Query - If multicast routing is not supported on other switches in your network, you can use IGMP Snooping and IGMP Query (page 234) to monitor IGMP service requests passing between multicast clients and servers, and dynamically configure the switch ports which need to forward multicast traffic. IGMP Snooping conserves bandwidth on network segments where no node has expressed interest in receiving a specific multicast service. For switches that do not support multicast routing, or where multicast routing is already enabled on other switches in the local network segment, IGMP Snooping is the only service required to support multicast filtering.

When using IGMPv3 snooping, service requests from IGMP Version 1, 2 or 3 hosts are all forwarded to the upstream router as IGMPv3 reports. The primary enhancement provided by IGMPv3 snooping is in keeping track of information about the specific multicast sources which downstream IGMPv3 hosts have requested or refused. The switch maintains information about both multicast groups and channels, where a group indicates a multicast flow for which the hosts have not requested a specific source (the only option for IGMPv1 and v2 hosts unless statically configured on the switch), and a channel indicates a flow for which the hosts have requested service from a specific source. For IGMPv1/v2 hosts, the source address of a channel is always null (indicating that any source is acceptable), but for IGMPv3 hosts, it may include a specific address when requested.

Only IGMPv3 hosts can request service from a specific multicast source. When downstream hosts request service from a specific source for a multicast service, these sources are all placed in the Include list, and traffic is forwarded to the hosts from each of these sources. IGMPv3 hosts may also request that service be forwarded from any source except for those specified. In this case, traffic is filtered from sources in the Exclude list, and forwarded from all other available sources.

NOTE: When the switch is configured to use IGMPv3 snooping, the snooping version may be downgraded to version 2 or version 1, depending on the version of the IGMP query packets detected on each VLAN.

NOTE: IGMP snooping will not function unless a multicast router port is enabled on the switch. This can be accomplished in one of two ways. A static router port can be manually configured (see the *Specifying Static Interfaces for a Multicast Router* section on page 236.) Using this method, the router port is never timed out, and will continue to function until explicitly removed. The other method relies on the switch to dynamically create multicast routing ports whenever multicast routing protocol packets or IGMP query packets are detected on a port.

NOTE: A maximum of up to 1023 multicast entries can be maintained for IGMP snooping. Once the table is full, no new entries are learned. Any subsequent multicast traffic not found in the table is dropped if unregistered-flooding is disabled (default behavior) and no router port is configured in the attached VLAN, or flooded throughout the VLAN if unregistered-flooding is enabled (see the *Configuring IGMP Snooping and Query Parameters* section on page 234).

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 236). 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 238).

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.

The following table lists the options on this page:

Multicast - IGMP Snooping (General) Options	
IGMP Snooping Status	<p>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)</p> <p>When IGMP snooping is enabled globally, the per VLAN interface settings for IGMP snooping take precedence (see the <i>Setting IGMP Snooping Status per Interface</i> section on page 239).</p> <p>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.</p>
Proxy Reporting Status	<p>Enables IGMP Snooping with Proxy Reporting. (Default: Disabled)</p> <p>When proxy reporting is enabled with this command, the switch performs <i>IGMP Snooping with Proxy Reporting</i> (as defined in DSL Forum TR-101, April 2006), including last leave, and query suppression.</p> <p>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.</p> <p>When proxy reporting is disabled, all IGMP reports received by the switch are forwarded natively to the upstream multicast routers.</p>
TCN Flood	<p>Enables flooding of multicast traffic if a spanning tree topology change notification (TCN) occurs. (Default: Disabled)</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>The proxy query and unsolicited MRD request are flooded to all VLAN ports except for the receiving port when the switch receives such packets.</p>
TCN Query Solicit	<p>Sends out an IGMP general query solicitation when a spanning tree topology change notification (TCN) occurs. (Default: Disabled)</p> <p>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.</p>
Router Alert Option	<p>Discards any IGMPv2/v3 packets that do not include the Router Alert option. (Default: Disabled)</p> <p>As described in Section 9.1 of RFC 3376 for IGMP Version 3, the Router Alert Option can be used to protect against DOS attacks. One common method of attack is launched by an intruder who takes over the role of querier, and starts overloading multicast hosts by sending a large number of group-and-source- specific queries, each with a large source list and the Maximum Response Time set to a large value.</p> <p>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.</p>
Unregistered Data Flooding	<p>Floods unregistered multicast traffic into the attached VLAN. (Default: Disabled)</p> <p>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.</p>

Multicast - IGMP Snooping (General) Options	
Forwarding Priority	Assigns a CoS priority to all multicast traffic. (Range: 0-7, where 7 is the highest priority; Default: Disabled) This parameter can be used to set a high priority for low-latency multicast traffic such as a video-conference, or to set a low priority for normal multicast traffic not sensitive to latency.
Version Exclusive	Discards any received IGMP messages which use a version different to that currently configured by the IGMP Version attribute. (Default: Disabled)
IGMP Unsolicited Report Interval	Specifies how often the upstream interface should transmit unsolicited IGMP reports when proxy reporting is enabled. (Range: 1-65535 seconds, Default: 400 seconds) When a new upstream interface (that is, uplink port) starts up, the switch sends unsolicited reports for all currently learned multicast channels via the new upstream interface. This command only applies when proxy reporting is enabled.
Router Port Expire Time	The time the switch waits after the previous querier stops before it considers it to have expired. (Range: 1-65535, Recommended Range: 300-500 seconds, Default: 300)
IGMP Snooping Version	Sets the protocol version for compatibility with other devices on the network. This is the IGMP Version the switch uses to send snooping reports. (Range: 1-3; Default: 2) This attribute configures the IGMP report/query version used by IGMP snooping. Versions 1 - 3 are all supported, and versions 2 and 3 are backward compatible, so the switch can operate with other devices, regardless of the snooping version employed.
Querier Status	When enabled, the switch can serve as the Querier, which is responsible for asking hosts if they want to receive multicast traffic. This feature is not supported for IGMPv3 snooping. (Default: Disabled)

Perform these steps 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**.

FIG. 276 Configuring General Settings for IGMP Snooping

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 the *Configuring IGMP Snooping and Query Parameters* section on page 234) before a multicast router port can take effect.

The following table lists the options on this page:

Multicast - IGMP Snooping (Multicast Router) Options	
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.

Multicast - IGMP Snooping (Multicast Router) Options	
Port or Trunk	Specifies the interface attached to a multicast router.
Show Static Multicast Router	
VLAN	Selects the VLAN for which to display any configured static multicast routers.
Interface	Shows the interface to which the specified static multicast routers are attached.
Show Current Multicast Router	
VLAN	Selects the VLAN for which to display any currently active multicast routers.
Interface	Shows the interface to which an active multicast router is attached.
Type	Shows if this entry is static or dynamic.
Expire	Time until this dynamic entry expires.

Perform these steps 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**.

Multicast > IGMP Snooping > Multicast Router

Action: Add Static Multicast Router

VLAN: 1

Interface: Port 1 Trunk

Apply Revert

FIG. 277 Configuring a Static Interface for a Multicast Router

Perform these steps 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.

Multicast > IGMP Snooping > Multicast Router

Action: Show Static Multicast Router

VLAN: 1

Static Multicast Router Interface List Total: 6

<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

Delete Revert

FIG. 278 Showing Static Interfaces Attached to a Multicast Router

Multicast routers that are attached to ports on the switch use information obtained from IGMP, along with a multicast routing protocol (such as PIM) to support IP multicasting across the Internet. These routers may be dynamically discovered by the switch or statically assigned to an interface on the switch.

Perform these steps 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.

The screenshot shows the 'Multicast > IGMP Snooping > Multicast Router' page. The 'Action' dropdown is set to 'Show Current Multicast Router'. The 'VLAN' dropdown is set to '1'. Below this, a table titled 'Multicast Router Interface Information' shows a total of 2 interfaces:

Interface	Type	Expire
Eth 1 / 1	Static	
Eth 1 / 4	Dynamic	0:4:59

FIG. 279 Showing Current Interfaces Attached to a Multicast Router

Assigning Interfaces to Multicast Services

Use the Multicast > IGMP Snooping > IGMP Member (Add Static Member) page to statically assign a multicast service to an interface.

Multicast filtering can be dynamically configured using IGMP Snooping and IGMP Query messages (see the *Configuring IGMP Snooping and Query Parameters* section on page 234). 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.

The following table lists the options on this page:

Multicast - IGMP Snooping (IGMP Member) Options	
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 router.
Multicast IP	The IP address for a specific multicast service.

Perform these steps 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**.

The screenshot shows the 'Multicast > IGMP Snooping > IGMP Member' page. The 'Action' dropdown is set to 'Add Static Member'. The 'VLAN' dropdown is set to '1'. The 'Interface' section has 'Port' selected with a dropdown set to '1', and 'Trunk' is also set to '1'. The 'Multicast IP' text box contains '224.1.1.1'. At the bottom, there are 'Apply' and 'Revert' buttons.

FIG. 280 Assigning an Interface to a Multicast Service

Perform these steps 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.

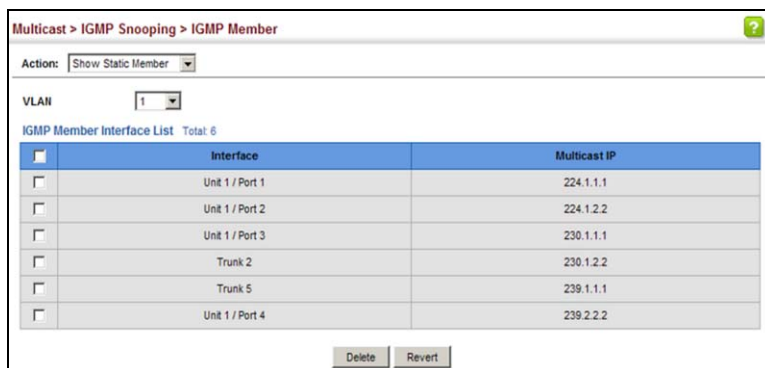


FIG. 281 Showing Static Interfaces Assigned to a Multicast Service

Setting IGMP Snooping Status per Interface

Use the Multicast > IGMP Snooping > Interface (Configure VLAN) page to configure IGMP snooping attributes for a VLAN. To configure snooping globally, refer to the *Configuring IGMP Snooping and Query Parameters* section on page 234.

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.*

The following table lists the options on this page:

Multicast - IGMP Snooping (Interface) Options	
VLAN	ID of configured VLANs. (Range: 1-4094)
IGMP Snooping Status	<p>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)</p> <p>When IGMP snooping is enabled globally (see page 234), 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.</p>
Version Exclusive	<p>Discards any received IGMP messages (except for multicast protocol packets) which use a version different to that currently configured by the IGMP Version attribute. (Options: Enabled, Using Global Status; Default: Using Global Status)</p> <p>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.</p>
Immediate Leave Status	<p>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)</p> <p>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.</p> <p>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.</p> <p>This attribute is only effective if IGMP snooping is enabled, and IGMPv2 or IGMPv3 snooping is used.</p> <p>If immediate leave is enabled, the following options are provided:</p> <ul style="list-style-type: none"> • By Group - The switch assumes that only one host is connected to the interface. Therefore, immediate leave should only be enabled on an interface if it is connected to only one IGMP-enabled device, either a service host or a neighbor running IGMP snooping. • By Host IP - The switch will not send out a group-specific query when an IGMPv2/v3 leave message is received. But will check if there are other hosts joining the multicast group. Only when all hosts on that port leave the group will the member port be deleted.
Multicast Router Discovery	MRD is used to discover which interfaces are attached to multicast routers. (Default: Disabled)
General Query Suppression	<p>Suppresses general queries except for ports attached to downstream multicast hosts. (Default: Disabled)</p> <p>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.</p>
Proxy Reporting	<p>Enables IGMP Snooping with Proxy Reporting. (Options: Enabled, Disabled, Using Global Status; Default: Using Global Status)</p> <p>When proxy reporting is enabled with this command, the switch performs <i>IGMP Snooping with Proxy Reporting</i> (as defined in DSL Forum TR-101, April 2006), including last leave, and query suppression.</p> <p>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.</p> <p>Rules Used for Proxy Reporting</p> <p>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.</p> <p>When IGMP Proxy Reporting is enabled, the source address is based on the following criteria:</p> <ul style="list-style-type: none"> • 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	<p>Sets the protocol version for compatibility with other devices on the network. This is the IGMP Version the switch uses to send snooping reports. (Options: 1-3, Using Global Version; Default: Using Global Version)</p> <p>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.</p>

Multicast - IGMP Snooping (Interface) Options	
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 234), or as a proxy host when IGMP snooping proxy reporting is enabled (page 234).
Query Response Interval	The maximum time the system waits for a response to general queries. (Range: 10-31740 tenths of a second in multiples of 10; Default: 10 seconds) This attribute applies when the switch is serving as the querier (page 234), or as a proxy host when IGMP snooping proxy reporting is enabled (page 234).
Last Member Query Interval	The interval to wait for a response to a group-specific or group-and-source-specific query message. (Range: 1-31744 tenths of a second in multiples of 10; Default: 1 second) When a multicast host leaves a group, it sends an IGMP leave message. When the leave message is received by the switch, it checks to see if this host is the last to leave the group by sending out an IGMP group-specific or group-and-source-specific query message, and starts a timer. If no reports are received before the timer expires, the group record is deleted, and a report is sent to the upstream multicast router. A reduced value will result in reduced time to detect the loss of the last member of a group or source, but may generate more burst traffic. This attribute will take effect only if IGMP snooping proxy reporting is enabled (page 234) or IGMP querier is enabled (page 234).
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).

Perform these steps 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**.

Multicast > IGMP Snooping > Interface

Action: Configure VLAN

VLAN: 1

IGMP Snooping Status: Enabled

Version Exclusive: Using Global Status

Immediate Leave Status: Enabled By-Group

Multicast Router Discovery: Enabled

General Query Suppression: Enabled

Proxy Reporting: Using Global Status

Interface Version: Using Global Version

Query Interval (2-31744): seconds

Query Response Interval (10-31740): (1/10 seconds, multiple of 10)

Last Member Query Interval (1-31744): (1/10 seconds, multiple of 10)

Last Member Query Count (1-255):

Proxy (Query) Address:

Apply Revert

FIG. 282 Configuring IGMP Snooping on a VLAN

Perform these steps to show the interface settings for IGMP snooping:

1. Click **Multicast > IGMP Snooping > Interface**.
2. Select **Show VLAN Information** from the Action list.

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

FIG. 283 Showing Interface Settings for IGMP Snooping

Filtering IGMP Query Packets and Multicast Data

Use the Multicast > IGMP Snooping > Interface (Configure Interface) page to configure an interface to drop IGMP query packets or multicast data packets.

The following table lists the options on this page:

Multicast - IGMP Snooping (Interface) Options	
Interface	Port or Trunk identifier
IGMP Query Drop	Configures an interface to drop any IGMP query packets received on the specified interface. If this switch is acting as a Querier, this prevents it from being affected by messages received from another Querier.
Multicast Data Drop	Configures an interface to stop multicast services from being forwarded to users attached to the downstream port (i.e., the interfaces specified by this command).

Perform these steps to drop IGMP query packets or multicast data packets:

1. Click **Multicast > IGMP Snooping > Interface**.
2. Select **Configure Interface** from the Action list.
3. Select **Port** or **Trunk** interface.
4. Enable the required drop functions for any interface.
5. Click **Apply**.

Port	IGMP Query Drop	Multicast Data Drop
1	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled
2	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled
3	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled
4	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled
5	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled

FIG. 284 Dropping IGMP Query or Multicast Data 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 234).

The following table lists the options on this page:

Multicast - IGMP Snooping (Forwarding Entry) Options	
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.

To show multicast groups learned through IGMP snooping, click **Multicast > IGMP Snooping > Forwarding Entry**.

Multicast > IGMP Snooping > Forwarding Entry						
IGMP Snooping Forwarding Entry List Total: 10						
VLAN	Group Address	Source Address	Interface	Up Time	Expire	Count
1	224.1.1.1	*	Eth 1 / 9 (Router Port)	00:00:06:46		2 (Port)
			Eth 1 / 11 (Member Port)	00:00:06:46	03:46	1 (Host)
1	224.1.1.2	192.168.1.2	Eth 1 / 9 (Router Port)		02:24	1 (Port)
2	224.1.1.3	*	Eth 1 / 9 (Router Port)	00:00:16:14		1 (Port)
2	239.255.255.250	*	Eth 1 / 9 (Router Port)	00:00:06:47		2 (Port)
			Eth 1 / 11 (Member Port)	00:00:06:47	03:46	1 (Host)

Clear Click this button to clear all IGMP Snooping dynamic groups.

FIG. 285 Showing Multicast Groups Learned by IGMP Snooping

Displaying IGMP Snooping Statistics

Use the Multicast > IGMP Snooping > Statistics pages to display IGMP snooping protocol-related statistics for the specified interface.

The following table lists the options on this page:

Multicast - IGMP Snooping (Statistics) Options	
VLAN	VLAN identifier. (Range: 1-4094)
Port	Port identifier. (Range: 1-10/26/28/52)
Trunk	Trunk identifier. (Range: 1-8)
Query Statistics	
Other Querier	IP address of remote querier on this interface
Other Querier Expire	Time after which remote querier is assumed to have expired.
Other Querier Uptime	Time remote querier has been up.
Self Querier	IP address of local querier on this interface.
Self Querier Expire	Time after which local querier is assumed to have expired.
Self Querier Uptime	Time local querier has been up.
General Query Received	The number of general queries received on this interface.
General Query Sent	The number of general queries sent from this interface.
Specific Query Received	The number of specific queries received on this interface.
Specific Query Sent	The number of specific queries sent from this interface.
Warn Rate Limit	The rate at which received query messages of the wrong version type cause the Vx warning count to increment. Note that "0 sec" means that the Vx warning count is incremented for each wrong message version received.
V1 Warning Count	The number of times the query version received (Version 1) does not match the version configured for this interface.
V2 Warning Count	The number of times the query version received (Version 2) does not match the version configured for this interface.

Multicast - IGMP Snooping (Statistics) Options	
V3 Warning Count	The number of times the query version received (Version 3) does not match the version configured for this interface.
VLAN, Port, and Trunk Statistics	
<i>Input Statistics</i>	
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.
<i>Output Statistics</i>	
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.

Perform these steps 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.

Action: Show Query Statistics

VLAN: 1

Query Statistics

Other Querier	None
Other Querier Expire	00(m):00(s)
Other Querier Uptime	00(h):00(m):00(s)
Self Querier	192.168.1.1
Self Querier Expire	00(m):00(s)
Self Querier Uptime	00(h):00(m):00(s)
General Query Received	0
General Query Sent	0
Specific Query Received	0
Specific Query Sent	0
Warn Rate Limit	0 sec.
V1 Warning Count	0
V2 Warning Count	0
V3 Warning Count	0

Clear All Click this button to clear all IGMP Snooping statistics.

Refresh

FIG. 286 Displaying IGMP Snooping Statistics - Query

Perform these steps 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.

The screenshot shows the 'Multicast > IGMP Snooping > Statistics' page. The 'Action' dropdown is set to 'Show VLAN Statistics'. A 'VLAN' dropdown menu is set to '1'. The page displays two sections: 'Input Statistics' and 'Output Statistics'. Each section has a table with four columns: Report, Leave, G Query, and G(-S)-S Query. The 'Drop', 'Join Success', and 'Group' columns are also present. All values are 0. At the bottom, there are 'Clear' and 'Refresh' buttons.

Input Statistics			
Report	0	Drop	0
Leave	0	Join Success	0
G Query	0	Group	0
G(-S)-S Query	0		

Output Statistics			
Report	0	Drop	0
Leave	0	Group	0
G Query	0		
G(-S)-S Query	0		

FIG. 287 Displaying IGMP Snooping Statistics - VLAN

Perform these steps 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.

The screenshot shows the 'Multicast > IGMP Snooping > Statistics' page. The 'Action' dropdown is set to 'Show Port Statistics'. A 'Port' dropdown menu is set to '1'. The page displays two sections: 'Input Statistics' and 'Output Statistics'. Each section has a table with four columns: Report, Leave, G Query, and G(-S)-S Query. The 'Drop', 'Join Success', and 'Group' columns are also present. All values are 0. At the bottom, there is a 'Refresh' button.

Input Statistics			
Report	0	Drop	0
Leave	0	Join Success	0
G Query	0	Group	0
G(-S)-S Query	0		

Output Statistics			
Report	0		
Leave	0		
G Query	0		
G(-S)-S Query	0		

FIG. 288 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. The following table lists the options on this page:

Multicast - IGMP Snooping (Filter) Options	
IGMP Filter Status	Enables IGMP filtering and throttling globally for the switch. (Default: Disabled)

Perform these steps 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**.

FIG. 289 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.

The following table lists the options on this page:

Multicast - IGMP Snooping (Filter) Options	
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.

Perform these steps 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**.

Multicast > IGMP Snooping > Filter

Step: 2. Configure Profile Action: Add

Profile ID (1-4294967295) 19

Access Mode Permit

Apply Revert

FIG. 290 Creating an IGMP Filtering Profile

Perform these steps 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.

Multicast > IGMP Snooping > Filter

Step: 2. Configure Profile Action: Show

IGMP Snooping Filter Profile List Total: 1

Profile ID	Action Mode
19	Permit

Delete Revert

FIG. 291 Showing the IGMP Filtering Profiles Created

Perform these steps 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**.

Multicast > IGMP Snooping > Filter

Step: 2. Configure Profile Action: Add Multicast Group Range

Profile ID 19

Start Multicast IP Address 239.2.3.1

End Multicast IP Address 239.2.3.200

Apply Revert

FIG. 292 Adding Multicast Groups to an IGMP Filtering Profile

Perform these steps 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.




FIG. 293 Showing the Groups Assigned to an IGMP Filtering Profile

Configuring IGMP Filtering and Throttling for Interfaces

Use the Multicast > IGMP Snooping > Filter (Configure Interface) page to assign an IGMP filter profile to interfaces on the switch, or to throttle multicast traffic by limiting the maximum number of multicast groups an interface can join at the same time.

Command Usage

- IGMP throttling sets a maximum number of multicast groups that a port can join at the same time. When the maximum number of groups is reached on a port, the switch can take one of two actions; either deny or replace. If the action is set to deny, any new IGMP join reports will be dropped. If the action is set to replace, the switch randomly removes an existing group and replaces it with the new multicast group.

The following table lists the options on this page:

Multicast - IGMP Snooping (Filter) Options	
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: 1-511; Default: 511)
Current Multicast Groups	Displays the current multicast groups the interface has joined.
Throttling Actions Mode	Sets the action to take when the maximum number of multicast groups for the interface has been exceeded. (Default: Deny) <ul style="list-style-type: none"> • 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)

Perform these steps 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**.

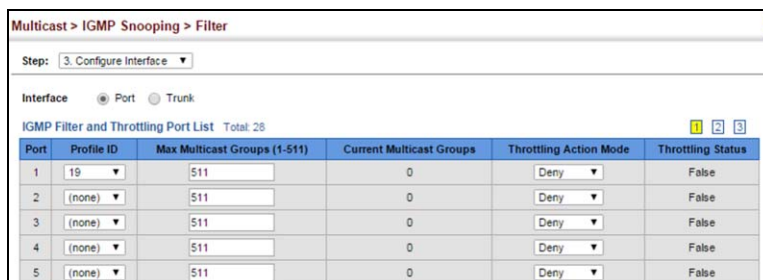


FIG. 294 Configuring IGMP Filtering and Throttling Interface Settings

MLD Snooping (Snooping and Query for IPv4)

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.

The following table lists the options on this page:

Multicast - MLD Snooping (General) Options	
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: <ul style="list-style-type: none"> 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.)

Perform these steps to configure general settings for MLD Snooping:

1. Click **Multicast > MLD Snooping > General**.
2. Adjust the settings as required.
3. Click **Apply**.

FIG. 295 Configuring General Settings for MLD Snooping

Setting Immediate Leave Status for MLD Snooping per Interface

Use the Multicast > MLD Snooping > Interface page to configure Immediate Leave status for a VLAN.

The following table lists the options on this page:

Multicast - MLD Snooping (Interface) Options	
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.

Perform these steps 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**.

FIG. 296 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 the *Configuring MLD Snooping and Query Parameters* section on page 249) before a multicast router port can take effect.

The following table lists the options on this page:

Multicast - MLD Snooping (Interface) Options	
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.

Perform these steps 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**.

Multicast > MLD Snooping > Multicast Router

Action: Add Static Multicast Router

VLAN: 1

Interface: Port 1 Trunk

Apply Revert

FIG. 297 Configuring a Static Interface for an IPv6 Multicast Router

Perform these steps 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.

Multicast > MLD Snooping > Multicast Router

Action: Show Static Multicast Router

VLAN: 1

Static Multicast Router Interface List Total: 2

	Interface
<input type="checkbox"/>	Unit 1 / Port 1
<input type="checkbox"/>	Unit 1 / Port 2

Delete Revert

FIG. 298 Showing Static Interfaces Attached an IPv6 Multicast Router

Perform these steps 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.

Multicast > MLD Snooping > Multicast Router

Action: Show Current Multicast Router

VLAN: 1

Multicast Router Interface Information Total: 4

Interface	Type
Unit 1 / Port 4	Static
Unit 1 / Port 5	Dynamic
Trunk 2	Dynamic
Trunk 3	Dynamic

FIG. 299 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 the *Configuring MLD Snooping and Query Parameters* section on page 249). 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.

The following table lists the options on this page:

Multicast - MLD Snooping (MLD Member) Options	
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).

Perform these steps 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**.

FIG. 300 Assigning an Interface to an IPv6 Multicast Service

Perform these steps 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.

	Multicast IPv6 Address	Interface
<input type="checkbox"/>	FF02:01:01:01:01	Unit 1 / Port 1
<input type="checkbox"/>	FF02:01:01:01:02	Unit 1 / Port 2
<input type="checkbox"/>	FF01:1	Unit 1 / Port 12
<input type="checkbox"/>	FF01:2	Unit 1 / Port 13
<input type="checkbox"/>	FF01:3	Unit 1 / Port 14
<input type="checkbox"/>	FF01:4	Unit 1 / Port 15
<input type="checkbox"/>	FF01:5	Unit 1 / Port 16
<input type="checkbox"/>	FF02:01:01:01:FF	Trunk 3

FIG. 301 Showing Static Interfaces Assigned to an IPv6 Multicast Service

Perform these steps 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.

Multicast IPv6 Address	Interface	Type
FF02::01:01:01:01	Unit 1 / Port 1	User
FF02::01:01:01:02	Unit 1 / Port 2	User
FF01::1	Unit 1 / Port 12	User
FF11::2	Unit 1 / Port 13	Multicast Data
FF11::3	Unit 1 / Port 14	User
FF11::4	Unit 1 / Port 15	User
FF11::5	Unit 1 / Port 16	User
FF02::01:01:01:FF	Trunk 3	User
FF03::01:01:01:FF	Trunk 5	MLD Snooping

FIG. 302 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.

The following table lists the options on this page:

Multicast - MLD Snooping (Group Information) Options	
VLAN	VLAN identifier. (Range: 1-4094)
Interface	Port or trunk identifier
Group Address	The IP address for a specific multicast service
Type	The means by which each group was learned - MLD Snooping or Multicast Data.
Filter Mode	The filter mode is used to summarize the total listening state of a multicast address to a minimum set such that all nodes' listening states are respected. In Include mode, the router only uses the request list, indicating that the reception of packets sent to the specified multicast address is requested only from those IP source addresses listed in the hosts' source-list. In Exclude mode, the router uses both the request list and exclude list, indicating that the reception of packets sent to the given multicast address is requested from all IP source addresses, except for those listed in the exclude source-list and for any other sources where the source timer status has expired.
Filter Timer Elapse	The Filter timer is only used when a specific multicast address is in Exclude mode. It represents the time for the multicast address filter mode to expire and change to Include mode.
Request List	Sources included on the router's request list
Exclude List	Sources included on the router's exclude list

Perform these steps 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.

Request List Total: 3	
IPv6 Address	
:::01:02:03:01	
:::01:02:03:02	
:::01:02:03:03	

Exclude List Total: 3	
IPv6 Address	
:::02:02:03:01	
:::02:02:03:02	
:::02:02:03:03	

FIG. 303 Showing IPv6 Multicast Services and Corresponding Sources

IP Tools

This chapter provides information on network functions including:

- **Ping** - Sends ping message to another node on the network.
- **Trace Route** - Sends ICMP echo request packets to another node on the network.
- **Address Resolution Protocol** - Describes how to configure ARP aging time, proxy ARP, or static addresses. Also shows how to display dynamic entries in the ARP cache.

Using the Ping Function

Use the Tool > Ping page to send ICMP echo request packets to another node on the network.

The following table lists the options on this page:

Tools - Ping Options	
Host Name/IP Address	Alias or IPv4/IPv6 address of the host
Probe Count	Number of packets to send (Range: 1-16)
Packet Size	Number of bytes in a packet. (Range: 32-512 bytes for IPv4, 0-1500 bytes for IPv6) The actual packet size will be eight bytes larger than the size specified because the switch adds header information.

Command Usage

- Use the ping command to see if another site on the network can be reached.
- The following are some results of the ping command:
 - Normal response - The normal response occurs in one to ten seconds, depending on network traffic.
 - Destination does not respond - If the host does not respond, a "timeout" appears in ten seconds.
 - Destination unreachable - The gateway for this destination indicates that the destination is unreachable.
 - Network or host unreachable - The gateway found no corresponding entry in the route table.
- The same link-local address may be used by different interfaces/nodes in different zones (RFC 4007). Therefore, when specifying a link-local address, include zone-id information indicating the VLAN identifier after the % delimiter. For example, FE80::7272%1 identifies VLAN 1 as the interface.

Perform these steps to ping another device on the network:

1. Click **Tool > Ping**.
2. Specify the target device and ping parameters.
3. Click **Apply**.

Tool > Ping

Host Name/IP Address:

Probe Count (1-16):

Data Size (IPv4: 32-512, IPv6: 0-1500): bytes

Note: For IPv4 Data Size,
 0 - 31 changed to 32 bytes
 32 - 512 is valid input
 513 - 1500 changed to 512 bytes
 < 0 or > 1500 not valid input

Result

```

PING to 192.168.2.99, by 5 of 32-byte payload ICMP packets, timeout is 3 seconds
response time: 0 ms
response time: 0 ms
response time: 0 ms
response time: 0 ms
response time: 0 ms

Ping statistics for 192.168.2.99:
 5 packets transmitted, 5 packets received (100%), 0 packets lost (0%)
Approximate round trip times:
  Minimum = 0 ms, Maximum = 0 ms, Average = 0 ms
  
```

FIG. 304 Pinging a Network Device

Using the Trace Route Function

Use the Tool > Trace Route page to show the route packets take to the specified destination.

The following table lists the options on this page:

Tools - Ping Options	
Destination Address	Alias or IPv4/IPv6 address of the host
IPv4 Max Failures	The maximum number of failures before which the trace route is terminated. (Fixed: 5)
IPv6 Max Failures	The maximum number of failures before which the trace route is terminated. (Range: 1-255; Default: 5)

Command Usage

- Use the trace route function to determine the path taken to reach a specified destination.
- A trace terminates when the destination responds, when the maximum timeout (TTL) is exceeded, or the maximum number of hops is exceeded.
- The trace route function first sends probe datagrams with the TTL value set at one. This causes the first router to discard the datagram and return an error message. The trace function then sends several probe messages at each subsequent TTL level and displays the round-trip time for each message. Not all devices respond correctly to probes by returning an *ICMP port unreachable* message. If the timer goes off before a response is returned, the trace function prints a series of asterisks and the *Request Timed Out* message. A long sequence of these messages, terminating only when the maximum timeout has been reached, may indicate this problem with the target device.
- The same link-local address may be used by different interfaces/nodes in different zones (RFC 4007). Therefore, when specifying a link-local address, include zone-id information indicating the VLAN identifier after the % delimiter. For example, FE80::7272%1 identifies VLAN 1 as the interface from which the trace route is sent.

Perform these steps to trace the route to another device on the network:

1. Click **Tool > Trace Route**.
2. Specify the target device.
3. Click **Apply**.

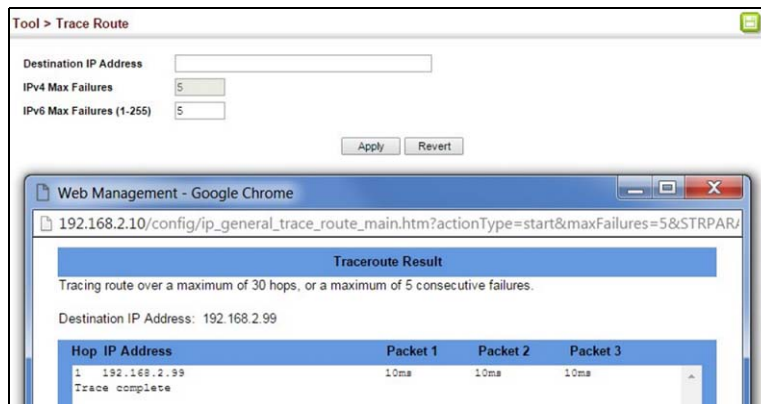


FIG. 305 Tracing the Route to a Network Device

Address Resolution Protocol

If IP routing is enabled, the router uses its routing tables to make routing decisions, and uses Address Resolution Protocol (ARP) to forward traffic from one hop to the next. ARP is used to map an IP address to a physical layer (i.e., MAC) address. When an IP frame is received by this router (or any standards-based router), it first looks up the MAC address corresponding to the destination IP address in the ARP cache. If the address is found, the router writes the MAC address into the appropriate field in the frame header, and forwards the frame on to the next hop. IP traffic passes along the path to its final destination in this way, with each routing device mapping the destination IP address to the MAC address of the next hop toward the recipient, until the packet is delivered to the final destination.

If there is no entry for an IP address in the ARP cache, the router will broadcast an ARP request packet to all devices on the network. The ARP request contains the following fields similar to that shown in this example:

Address Resolution Protocol	
destination IP address	10.1.0.19
destination MAC address	?
source IP address	10.1.0.253
source MAC address	00-00-ab-cd-00-00

When devices receive this request, they discard it if their address does not match the destination IP address in the message. However, if it does match, they write their own hardware address into the destination MAC address field and send the message back to the source hardware address. When the source device receives a reply, it writes the destination IP address and corresponding MAC address into its cache, and forwards the IP traffic on to the next hop. As long as this entry has not timed out, the router will be able to 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.

Displaying Dynamic or Local ARP Entries

Use the Tool > ARP page to display dynamic or local entries in the ARP cache. The ARP cache contains static entries, and entries for local interfaces, including subnet, host, and broadcast addresses. However, most entries will be dynamically learned through replies to broadcast messages.

Perform these steps to display all dynamic and local entries in the ARP cache:

1. Click **Tool > ARP**.
2. Select **Show Information** from the Step List.
3. Click **ARP Addresses**.

Tool > ARP			
ARP Address List Total: 1			
IP Address	MAC Address	Type	Interface
192.168.2.99	00-E0-4C-68-12-66	dynamic	VLAN 1

Clear Dynamic ARP

FIG. 306 Displaying ARP Entries

IP Services

This chapter describes the following IP services:

- **DNS** - Configures default domain names, identifies servers to use for dynamic lookup, and shows how to configure static entries.
- **DHCP Client** - Specifies the DHCP client identifier for an interface.
- **DHCP Relay** - Enables DHCP relay service for attached host devices, including DHCP option 82 information, and defines the servers to which client requests are forwarded.
- **DHCP Dynamic Provision** - Enables dynamic provision via DHCP.

NOTE: For information on DHCP snooping which is included in this folder, see the *DHCP Snooping* section on page 177.

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 the *Configuring a List of Name Servers* section on page 258).
- If one or more name servers are configured, but DNS is not yet enabled and the switch receives a DHCP packet containing a DNS field with a list of DNS servers, then the switch will automatically enable DNS host name-to-address translation.

The following table lists the options on this page:

IP Service - DNS Options	
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)

Perform these steps 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**.

FIG. 307 Configuring General Settings for DNS

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 the *Configuring General DNS Service Parameters* section on page 257). 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 the *Configuring a List of Name Servers* section on page 258).
- If all name servers are deleted, DNS will automatically be disabled.

The following table lists the options on this page:

IP Service - DNS Options	
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)

Perform these steps 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**.

FIG. 308 Configuring a List of Domain Names for DNS

Perform these steps to show the list domain names:

1. Click **IP Service > DNS**.
2. Select **Show Domain Names** from the Action list.

FIG. 309 Showing the List of Domain Names for DNS

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 the *Configuring General DNS Service Parameters* section on page 257).
- 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.

The following table lists the options on this page:

IP Service - DNS Options	
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.

Perform these steps 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**.

FIG. 310 Configuring a List of Name Servers for DNS

Perform these steps to show the list name servers:

1. Click **IP Service > DNS**.
2. Select **Show Name Servers** from the Action list.

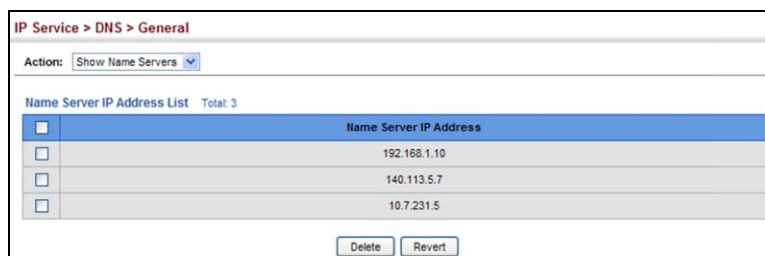


FIG. 311 Showing the List of Name Servers for DNS

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.

The following table lists the options on this page:

IP Service - DNS (Static Host Table) Options	
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.

Perform these steps 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**.



FIG. 312 Configuring Static Entries in the DNS Table

Perform these steps to show static entries in the DNS table:

1. Click **IP Service > DNS > Static Host Table**.
2. Select **Show** from the Action list.




FIG. 313 Showing Static Entries in the DNS Table

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.

The following table lists the options on this page:

IP Service - DNS (Cache) Options	
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.

To display entries in the DNS cache, click **IP Service > 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.

FIG. 314 Showing Entries in the DNS Cache

Dynamic Host Configuration Protocol

Dynamic Host Configuration Protocol (DHCP) can dynamically allocate an

IP address and other configuration information to network clients when they boot up. If a subnet does not already include a BOOTP or DHCP server, you can relay DHCP client requests to a DHCP server on another subnet, or configure the DHCP server on this switch to support that subnet.

When configuring the DHCP server on this switch, you can configure an address pool for each unique IP interface, or manually assign a static IP address to clients based on their hardware address or client identifier. The DHCP server can provide the host's IP address, domain name, gateway router and DNS server, information about the host's boot image including the TFTP server to access for download and the name of the boot file, or boot information for NetBIOS Windows Internet Naming Service (WINS).

Specifying a DHCP Client Identifier

Use the IP Service > DHCP > Client page to specify the DHCP client identifier for a VLAN interface.

Command Usage

- The class identifier is used identify the vendor class and configuration of the switch to the DHCP server, which then uses this information to decide on how to service the client or the type of information to return.
- The general framework for this DHCP option is set out in RFC 2132 (Option 60). This information is used to convey configuration settings or other identification information about a client, but the specific string to use should be supplied by your service provider or network administrator. Options 60, 66 and 67 statements can be added to the server daemon's configuration file.

Options 60, 66, and 67 Statements		
Option	Keyword	Statement Parameter
60	vendor-class-identifier	a string indicating the vendor class identifier
66	tftp-server-name	a string indicating the tftp server name
67	bootfile-name	a string indicating the bootfile name

- By default, DHCP option 66/67 parameters are not carried in a DHCP server reply. To ask for a DHCP reply with option 66/67 information, the DHCP client request sent by this switch includes a *parameter request list* asking for this information. Besides, the client request also includes a *vendor class identifier* that allows the DHCP server to identify the device, and select the appropriate configuration file for download. This information is included in Option 55 and 124.

Options 55 and 124 Statements		
Option	Keyword	Statement Parameter
55	dhcp-parameter-request-list	a list of parameters, separated by ','
124	vendor-class-identifier	a string indicating the vendor class identifier

- The server should reply with the TFTP server name and boot file name.
- Note that the vendor class identifier can be formatted in either text or hexadecimal, but the format used by both the client and server must be the same.

The following table lists the options on this page:

IP Service - DHCP (Client) Options	
VLAN	ID of configured VLAN
Vendor Class ID	The following options are supported when the check box is marked to enable this feature: <ul style="list-style-type: none"> • Default - The default string is the model number. • Text - A text string (Range: 1-32 characters) • Hex - A hexadecimal value (Range: 1-64 characters)

Perform these steps 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**.

The screenshot shows the configuration interface for DHCP Client options. The 'VLAN' is set to '1'. Under 'Vendor Class ID', the checkbox is checked, and the format is set to 'Default'. The text field contains the value 'ECS2000-28P'. At the bottom, there are 'Apply' and 'Revert' buttons.

FIG. 315 Specifying a DHCP Client Identifier

Configuring DHCP Relay Service

Use the IP Service > DHCP > Relay 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.

If DHCP relay is enabled, and this switch sees a DHCP request broadcast, it inserts its own IP address into the request so that the DHCP server will know the subnet where the client is located. Then, the switch forwards the packet to the DHCP server. When the server receives the DHCP request, it allocates a free IP address for the DHCP client from its defined scope for the DHCP client's subnet, and sends a DHCP response back to the DHCP relay agent (i.e., this switch). This switch then passes the DHCP response received from the server to the client.

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.




FIG. 316 Layer 3 DHCP Relay Service

Command Usage

- You must specify the IP address for at least one active DHCP server. Otherwise, the switch's DHCP relay agent will not be able to forward client requests to a DHCP server. Up to five DHCP servers can be specified in order of preference. If any of the specified DHCP server addresses are not located in the same network segment with this switch, specify the default router through which this switch can reach other IP subnetworks (see the *Configuring the IPv4 Default Gateway* section on page 265 or the *Configuring the IPv6 Default Gateway* section on page 267).

- DHCP relay configuration will be disabled if an active DHCP server is detected on the same network segment.
- DHCP Snooping Information Option 82 (see page 177) 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.

The following table lists the options on this page:

IP Service - DHCP (Relay) Options	
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: <ul style="list-style-type: none"> • 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)
Server IP Address	Addresses of DHCP servers or relay servers to be used by the switch's DHCP relay agent in order of preference.
Restart DHCP Relay	Use this button to re-initialize DHCP relay service.

Perform these steps to configure DHCP relay service:

1. Click **IP Service > DHCP > Relay**.
2. Enable or disable Option 82 relay information.
3. Set the Option 82 policy to specify how to handle Option 82 information already contained in DHCP client request packets.
4. Select whether or not to include the use of sub-type and sub-length fields for the circuit-ID (CID) and remote-ID (RID) in Option 82 information generated by the switch.
5. Enter up to five IP addresses for DHCP servers or relay servers in order of preference for any VLAN.
6. Click **Apply**.

FIG. 317 Configuring DHCP Relay Service

Enabling DHCP Dynamic Provision

Use the IP Service > DHCP > Dynamic Provision to enable dynamic provisioning via DHCP.

Command Usage

DHCPD is the daemon used by Linux to dynamically configure TCP/IP information for client systems. To support DHCP option 66/67, you have to add corresponding statements to the configuration file of DHCPD. Information on how to complete this process are described in the *Downloading a Configuration File and Other Parameters Provided by a DHCP Server* section in the *CLI Reference Guide*.

Some alternative commands which can be added to the DHCPD to complete the dynamic provisioning process are also described under the `ip dhcp dynamic-` provision command in the *CLI Reference Guide*.

By default, the parameters for DHCP option 66/67 are not carried by the reply sent from the DHCP server. To ask for a DHCP reply with option 66/67, the client can inform the server that it is interested in option 66/67 by sending a DHCP request that includes a 'parameter request list' option. Besides this, the client can also send a DHCP request that includes a 'vendor class identifier' option to the server so that the DHCP server can identify the device, and determine what information should be given to requesting device.

The following table lists the options on this page:

IP Service - DHCP (Dynamic Provision) Options	
Dynamic Provision via DHCP Status	Enables dynamic provisioning via DHCP. (Default: Disabled)

Perform these steps to enable dynamic provisioning via DHCP:

1. Click **IP Service > DHCP > Dynamic Provision**.
2. Mark the Enable box if dynamic provisioning is configured on the DHCP daemon, and required for boot-up.
3. Click **Apply**.




FIG. 318 Enabling Dynamic Provisioning via DHCP

IP Configuration

This chapter describes how to configure an IP interface for management access to the switch over the network. This switch supports both IP Version 4 and Version 6, and can be managed simultaneously through either of these address types. You can manually configure a specific IPv4 or IPv6 address, or direct the switch to obtain an IPv4 address from a BOOTP or DHCP server. An IPv6 address can either be manually configured or dynamically generated.

This chapter provides information on network functions including:

- **IPv4 Configuration** - Sets an IPv4 address for management access.
- **IPv6 Configuration** - Sets an IPv6 address for management access.

Setting the Switch's IP Address (IP Version 4)

This section describes how to configure an IPv4 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 IPv6 address, see the *Setting the Switch's IP Address (IP Version 6)* section on page 267.

Configuring the IPv4 Default Gateway

Use the System > IP (Configure Global) page to configure an IPv4 default gateway for the switch.

The following table lists the options on this page:

System - IP Options	
Gateway IP Address	IP address of the gateway router between the switch and management stations that exist on other network segments. (Default: None) An IP default gateway must be defined if the management station is located in a different IP segment. An IP default gateway can only be successfully set when a network interface that directly connects to the gateway has been configured on the switch.

Perform these steps to configure an IPv4 default gateway for the switch:

1. Click **System > IP**.
2. Select **Configure Global** from the Action list.
3. Enter the IPv4 default gateway.
4. Click **Apply**.

The screenshot shows the 'System > IP' configuration page. At the top, there is a 'Step:' dropdown menu set to '1. Configure Global'. Below this, the 'Gateway IP Address' field is a text input box containing the value '192.168.0.1'. At the bottom right of the form, there are two buttons: 'Apply' and 'Revert'.

FIG. 319 Configuring the IPv4 Default Gateway

Configuring IPv4 Interface Settings

Use the System > IP (Configure Interface - Add Address) page to configure an IPv4 address for the switch. The default IPv4 address for VLAN 1 is set to 192.168.2.10 using the subnet mask 255.255.255.0. To change the switch's default settings to values that are compatible with your network, you may need to establish a default gateway between the switch and management stations that exist on another network segment.

You can direct the device to obtain an address from a BOOTP or DHCP server, or manually configure a static IP address. Valid IP addresses consist of four decimal numbers, 0 to 255, separated by periods. Anything other than this format will not be accepted.

The following table lists the options on this page:

System - IP Options	
VLAN	ID of the VLAN 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; Default: VLAN 1)
IP Address Mode	Specifies whether IP functionality is enabled via manual configuration (User Specified), Dynamic Host Configuration Protocol (DHCP), or Boot Protocol (BOOTP). If DHCP/BOOTP is enabled, IP will not function until a reply has been received from the server. Requests will be broadcast periodically by the switch for an IP address. DHCP/BOOTP responses can include the IP address, subnet mask, and default gateway. (Default: User Specified)

System - IP Options	
IP Address Type	Specifies a primary or secondary IP address. An interface can have only one primary IP address, but can have many secondary IP addresses. In other words, secondary addresses need to be specified if more than one IP subnet can be accessed through this interface. For initial configuration, set this parameter to Primary. (Options: Primary, Secondary; Default: Primary) Note that a secondary address cannot be configured prior to setting the primary IP address, and the primary address cannot be removed if a secondary address is still present. Also, if any router or switch in a network segment uses a secondary address, all other routers/switches in that segment must also use a secondary address from the same network or subnet address space.
IP Address	IP Address of the VLAN. Valid IP addresses consist of four numbers, 0 to 255, separated by periods. (Default: 192.168.2.10)
Subnet Mask	This mask identifies the host address bits used for routing to specific subnets. (Default: 255.255.255.0)
Restart DHCP	Requests a new IP address from the DHCP server.

Perform these steps to set a static IPv4 address for the switch:

1. Click **System > IP**.
2. Select **Configure Interface** from the Step list.
3. Select **Add Address** from the Action list.
4. Select any configured VLAN, set the IP Address Mode to **User Specified**, set the IP Address Type to **Primary** if no address has yet been configured for this interface, and then enter the IP address and subnet mask.
5. Select **Primary** or **Secondary** Address Type.
6. Click **Apply**.

The screenshot shows the 'System > IP' configuration page. At the top, 'Step: 1. Configure Interface' and 'Action: Add Address' are selected. The 'VLAN' is set to '1'. 'IP Address Mode' is 'User Specified', and 'IP Address Type' is 'Primary'. The 'IP Address' field contains '192.168.0.2' and the 'Subnet Mask' field contains '255.255.255.0'. There is a 'Restart DHCP' button with a tooltip that says 'Click this button to resend DHCP client request.' At the bottom, there are 'Apply' and 'Revert' buttons.

FIG. 320 Configuring a Static IPv4 Address

Perform these steps to obtain an dynamic IPv4 address through DHCP/BOOTP for the switch:

1. Click **System > IP**.
2. Select **Configure Interface** from the Step list.
3. Select **Add Address** from the Action list.
4. Select any configured VLAN, and set IP Address Mode to BOOTP or DHCP.
5. Click **Apply** to save your changes.
6. Click **Restart DHCP** to immediately request a new address. IP will be enabled but will not function until a BOOTP or DHCP reply is received. Requests are broadcast every few minutes using exponential back-off until IP configuration information is obtained from a BOOTP or DHCP server.

The screenshot shows the 'System > IP' configuration page. At the top, 'Step: 2. Configure Interface' and 'Action: Add Address' are selected. The 'VLAN' is set to '1'. 'IP Address Mode' is 'DHCP', and 'IP Address Type' is 'Primary'. The 'IP Address' and 'Subnet Mask' fields are empty. There is a 'Restart DHCP' button with a tooltip that says 'Click this button to resend DHCP client request.' At the bottom, there are 'Apply' and 'Revert' buttons.

FIG. 321 Configuring a Dynamic IPv4 Address

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.

Perform these steps to show the IPv4 address configured for an interface:

1. Click **System > IP**.
2. Select **Configure Interface** from the Step list.
3. Select **Show Address** from the Action list.
4. Select an entry from the VLAN list.

The screenshot shows the 'System > IP' configuration page. At the top, there are two dropdown menus: 'Step: 2. Configure Interface' and 'Action: Show Address'. Below these, there is a table with the following information:

VLAN	1
IP Address Mode	DHCP
IP Address	192.168.2.12

FIG. 322 Showing the Configured IPv4 Address for an Interface

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 the *Setting the Switch's IP Address (IP Version 4)* section on page 265.

Command Usage

- IPv6 includes two distinct address types - link-local unicast and global unicast. A link-local address makes the switch accessible over IPv6 for all devices attached to the same local subnet. Management traffic using this kind of address cannot be passed by any router outside of the subnet. A link-local address is easy to set up, and may be useful for simple networks or basic troubleshooting tasks. However, to connect to a larger network with multiple segments, the switch must be configured with a global unicast address. Both link-local and global unicast address types can either be dynamically assigned (using the Configure Interface page) or manually configured (using the Add IPv6 Address page).
- An IPv6 global unicast or link-local address can be manually configured (using the Add IPv6 Address page), or a link-local address can be dynamically generated (using the Configure Interface page).

Configuring the IPv6 Default Gateway

Use the System > IPv6 Configuration (Configure Global) page to configure an IPv6 default gateway for the switch.

The following table lists the options on this page:

System - IPv6 Configuration Options	
Default Gateway	<p>Sets the IPv6 address of the default next hop router to use when no routing information is known about an IPv6 address.</p> <ul style="list-style-type: none"> • An IPv6 default gateway can only be successfully set when a network interface that directly connects to the gateway has been configured on the switch. • An IPv6 address must be configured according to RFC 2373 <i>IPv6 Addressing Architecture</i>, 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.

Perform these steps to configure an IPv6 default gateway for the switch:

1. Click **System > IPv6 Configuration**.
2. Select **Configure Global** from the Action list.
3. Enter the IPv6 default gateway.
4. Click **Apply**.

The screenshot shows the 'System > IPv6 Configuration' page. At the top, there is a dropdown menu for 'Action: Configure Global'. Below this, there is a text input field for 'Default Gateway' containing the value '2001:db8:2222:7272::254'. At the bottom of the form, there are two buttons: 'Apply' and 'Revert'.

FIG. 323 Configuring the IPv6 Default Gateway

Configuring IPv6 Interface Settings

Use the System > 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 be configured with a link-local address. The switch's address auto-configuration function will automatically create a link-local address, as well as an IPv6 global address if router advertisements are detected on the local interface.
- The option to explicitly enable IPv6 creates a link-local address, but will not generate a global IPv6 address if auto-configuration is not enabled. In this case, you can manually configure a global unicast address (see the *Configuring an IPv6 Address* section on page 270).
- 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.

The following table lists the options on this page:

System - IPv6 Configuration Options	
VLAN	ID of a configured VLAN that 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 auto-configuration 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). <ul style="list-style-type: none"> • If a link local address has not yet been assigned to this interface, this command will dynamically generate one. The link-local address is made with an address prefix in the range of FE80~FEBF and a host portion based the switch's MAC address in modified EUI-64 format. It will also generate a global unicast address if a global prefix is included in received router advertisements. • When DHCPv6 is started, the switch may attempt to acquire an IP address prefix through stateful address auto-configuration. If router advertisements have the <i>other stateful configuration</i> flag set, the switch will attempt to acquire other non-address configuration information (such as a default gateway). • If auto-configuration is not selected, then an address must be manually configured using the Add IPv6 Address page described below.
Enable IPv6 Explicitly	Enables IPv6 on an interface and assigns it a link-local address. Note that when an explicit address is assigned to an interface, IPv6 is automatically enabled, and cannot be disabled until all assigned addresses have been removed. (Default: Disabled) Disabling this parameter does not disable IPv6 for an interface that has been explicitly configured with an IPv6 address.
MTU	Sets the size of the maximum transmission unit (MTU) for IPv6 packets sent on an interface. (Range: 1280-65535 bytes; Default: 1500 bytes) <ul style="list-style-type: none"> • The maximum value set in this field cannot exceed the MTU of the physical interface, which is currently fixed at 1500 bytes. • If a non-default value is configured, an MTU option is included in the router advertisements sent from this device. This option is provided to ensure that all nodes on a link use the same MTU value in cases where the link MTU is not otherwise well known. • IPv6 routers do not fragment IPv6 packets forwarded from other routers. However, traffic originating from an end-station connected to an IPv6 router may be fragmented. • All devices on the same physical medium must use the same MTU in order to operate correctly. • IPv6 must be enabled on an interface before the MTU can be set. If an IPv6 address has not been assigned to the switch, N/A is displayed in the MTU field.

System - IPv6 Configuration Options	
ND DAD Attempts	<p>The number of consecutive neighbor solicitation messages sent on an interface during duplicate address detection. (Range: 0-600, Default: 3)</p> <ul style="list-style-type: none"> 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 the <i>Configuring VLAN Groups</i> section on page 88). 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	<p>The interval between transmitting IPv6 neighbor solicitation messages on an interface. (Range: 1000-3600000 milliseconds)</p> <p>Default: 1000 milliseconds is used for neighbor discovery operations, 0 milliseconds is advertised in router advertisements.</p> <p>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.</p> <p>When a non-default value is configured, the specified interval is used both for router advertisements and by the router itself.</p>
ND Reachable-Time	<p>The amount of time that a remote IPv6 node is considered reachable after some reachability confirmation event has occurred. (Range: 0-3600000 milliseconds)</p> <p>Default: 30000 milliseconds is used for neighbor discovery operations, 0 milliseconds is advertised in router advertisements.</p> <ul style="list-style-type: none"> The time limit configured by this parameter allows the router to detect unavailable neighbors. During the neighbor discover process, an IPv6 node will multicast neighbor solicitation messages to search for neighbor nodes. For a neighbor node to be considered reachable, it must respond to the neighbor soliciting node with a neighbor advertisement message to become a confirmed neighbor, after which the reachable timer will be considered in effect for subsequent unicast IPv6 layer communications. This time limit is included in all router advertisements sent out through an interface, ensuring that nodes on the same link use the same time value. Setting the time limit to 0 means that the configured time is unspecified by this router.
Restart DHCPv6	<p>When DHCPv6 is restarted, the switch may attempt to acquire an IP address prefix through stateful address auto-configuration. If the router advertisements have the <i>other stateful configuration</i> flag set, the switch may also attempt to acquire other non-address configuration information (such as a default gateway) when DHCPv6 is restarted.</p> <p>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.</p> <ul style="list-style-type: none"> 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 auto-configuration, 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 auto-configuration, in which a DHCPv6 server does not assign stateful addresses to IPv6 hosts, but does assign stateless configuration settings.

Perform these steps to configure general IPv6 settings for the switch:

1. Click **System > IPv6 Configuration**.
2. Select **Configure Interface** from the Action list.
3. Specify the VLAN to configure.
4. Enable address auto-configuration, or enable IPv6 explicitly to automatically configure a link-local address and enable IPv6 on the selected interface. (To manually configure the link-local address, use the Add IPv6 Address page.) Set the MTU size, the maximum number of duplicate address detection messages, the neighbor solicitation message interval, and the amount of time that a remote IPv6 node is considered reachable.
5. Click **Apply**.

The screenshot shows the 'System > IPv6 Configuration' page. At the top, the 'Action' dropdown is set to 'Configure Interface'. Below this, the 'VLAN' is set to '1'. The 'Address Autoconfig' and 'Enable IPv6 Explicitly' options are both checked and labeled 'Enabled'. The 'MTU (1280-65535)' is set to '1500 bytes'. The 'ND DAD Attempts (0-600)' is set to '3'. The 'ND NS Interval (1000-3600000)' is set to '1000 ms'. The 'ND Reachable-Time (0-3600000)' is set to '30000 ms'. At the bottom, there is a 'Restart DHCPv6' button with a tooltip that says 'Click this button to restart DHCPv6 service.', and 'Apply' and 'Revert' buttons.

FIG. 324 Configuring General Settings for an IPv6 Interface

Configuring an IPv6 Address

Use the System > 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 the *Configuring IPv6 Interface Settings* section on page 268), will also automatically generate a link-local unicast address. The prefix length for a link-local address is fixed at 64 bits, and the host portion of the default address is based on the modified EUI-64 (Extended Universal Identifier) form of the interface identifier (i.e., the physical MAC address). Alternatively, you can manually configure the link-local address by entering the full address with a network prefix in the range of FE80~FEBF.
- To connect to a larger network with multiple subnets, you must configure a global unicast address. There are several alternatives to configuring this address type:
 - The global unicast address can be automatically configured by taking the network prefix from router advertisements observed on the local interface, and using the modified EUI-64 form of the interface identifier to automatically create the host portion of the address (see the *Configuring IPv6 Interface Settings* section on page 268).
 - 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.

The following table lists the options on this page:

System - IPv6 Configuration Options	
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	<p>Defines the address type configured for this interface.</p> <ul style="list-style-type: none"> 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. <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> <ul style="list-style-type: none"> 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. <p>The specified address replaces a link-local address that was automatically generated for the interface.</p>
IPv6 Address	IPv6 address assigned to this interface.

Perform these steps to configure an IPv6 address:

1. Click **System > 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**.

The screenshot shows the 'System > IPv6 Configuration' page. At the top, the 'Action' dropdown is set to 'Add IPv6 Address'. Below this, there are three main configuration fields: 'VLAN' is set to '1', 'Address Type' is set to 'Global', and 'IPv6 Address' is set to '2001:db8:2222:7272::72/96'. At the bottom of the form, there are two buttons: 'Apply' and 'Revert'.

FIG. 325 Configuring an IPv6 Address

Showing IPv6 Addresses

Use the System > IPv6 Configuration (Show IPv6 Address) page to display the IPv6 addresses assigned to an interface.

The following table lists the options on this page:

System - IPv6 Configuration Options	
VLAN	ID of a configured VLAN. By default, all ports on the switch are members of VLAN 1. However, the management station can be attached to a port belonging to any VLAN, as long as that VLAN has been assigned an IP address. (Range: 1-4094)
IPv6 Address Type	The address type (Global, EUI-64, Link Local)
IPv6 Address	<p>An IPv6 address assigned to this interface.</p> <p>In addition to the unicast addresses assigned to an interface, a node is also required to listen to the all-nodes multicast addresses FF01::1 (interface-local scope) and FF02::1 (link-local scope). FF01::1/16 is the transient interface-local multicast address for all attached IPv6 nodes, and FF02::1/16 is the link-local multicast address for all attached IPv6 nodes. The interface-local multicast address is only used for loopback transmission of multicast traffic. Link-local multicast addresses cover the same types as used by link-local unicast addresses, including all nodes (FF02::1), all routers (FF02::2), and solicited nodes (FF02::1:FFXX:XXXX) as described below.</p> <p>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.</p> <p>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.</p> <p>These additional addresses are displayed by the <i>show ip interface</i> command described in the CLI Reference Guide).</p>
Configuration Mode	Indicates if this address was automatically generated or manually configured.

Perform these steps to show the configured IPv6 addresses:

1. Click **System > IPv6 Configuration**.
2. Select Show IPv6 Address from the Action list.
3. Select a VLAN from the list.

The screenshot shows the 'System > IPv6 Configuration' page. The 'Action' dropdown is set to 'Show IPv6 Address'. The 'VLAN' dropdown is set to '1'. Below, the 'IPv6 Address List' shows two entries:

IPv6 Address Type	IPv6 Address	Configuration Mode
Global	2001:db8:2222:7272::72/96	Manual
Link Local	fe80::2e0:cff:fe00:fd%1/64	Auto

Buttons for 'Delete' and 'Revert' are visible at the bottom of the list.

FIG. 326 Showing Configured IPv6 Addresses

Showing the IPv6 Neighbor Cache

Use the System > IPv6 Configuration (Show IPv6 Neighbor Cache) page to display the IPv6 addresses detected for neighbor devices.

The following table lists the options on this page:

Show IPv6 Neighbors - Display Description	
IPv6 Address	IPv6 address of neighbor.
Age	The time since the address was verified as reachable (in seconds). A static entry is indicated by the value Permanent.
Link-layer Address	Physical layer MAC address.
State	<p>The following states are used for dynamic entries:</p> <ul style="list-style-type: none"> • Incomplete - Address resolution is being carried out on the entry. A neighbor solicitation message has been sent to the multicast address of the target, but it has not yet returned a neighbor advertisement message. • Invalid - An invalidated mapping. Setting the state to invalid dis-associates the interface identified with this entry from the indicated mapping (RFC 4293). • Reachable - Positive confirmation was received within the last ReachableTime interval that the forward path to the neighbor was functioning. While in Reachable state, the device takes no special action when sending packets. • Stale - More than the ReachableTime interval has elapsed since the last positive confirmation was received that the forward path was functioning. While in Stale state, the device takes no action until a packet is sent. • Delay - More than the ReachableTime interval has elapsed since the last positive confirmation was received that the forward path was functioning. A packet was sent within the last DELAY_FIRST_PROBE_TIME interval. If no reachability confirmation is received within this interval after entering the Delay state, the switch will send a neighbor solicitation message and change the state to Probe. • Probe - A reachability confirmation is actively sought by re-sending neighbor solicitation messages every RetransTimer interval until confirmation of reachability is received. • Unknown - Unknown state. <p>The following states are used for static entries:</p> <ul style="list-style-type: none"> • Incomplete - The interface for this entry is down. • Permanent - Indicates a static entry. • Reachable - The interface for this entry is up. Reachability detection is not applied to static entries in the IPv6 neighbor discovery cache.
VLAN	VLAN interface from which the address was reached.

Perform these steps to show neighboring IPv6 devices:

1. Click **System > IPv6 Configuration**.
2. Select **Show IPv6 Neighbors** from the Action list.

The screenshot shows the 'System > IPv6 Configuration' page. At the top, there is a dropdown menu for 'Action' with 'Show IPv6 Neighbor Cache' selected. Below this, a table titled 'Current Neighbor Cache Table' shows a total of 1 entry. The table has columns for IPv6 Address, Age, Link-layer Address, State, and VLAN. The entry shown is for IPv6 Address fe80::ce37:abff:fe5b:77bc, Age 26, Link-layer Address CC-37-AB-5B-77-BC, State Stale, and VLAN 1. A 'Clear' button is located below the table.

IPv6 Address	Age	Link-layer Address	State	VLAN
fe80::ce37:abff:fe5b:77bc	26	CC-37-AB-5B-77-BC	Stale	1

FIG. 327 Showing IPv6 Neighbors

Showing IPv6 Statistics

Use the System > 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 feedback 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.

The following table lists the options on this page:

Show IPv6 Statistics - Display 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 Errors	The number of input datagrams discarded because datagram frame didn't carry enough data.
Discards	The number of input IPv6 datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (e.g., for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting re-assembly.
Delivers	The total number of datagrams successfully delivered to IPv6 user- protocols (including ICMP). This counter is incremented at the interface to which these datagrams were addressed which might not be necessarily the input interface for some of the datagrams.
Reassembly Request Datagrams	The number of IPv6 fragments received which needed to be reassembled at this interface. Note that this counter is incremented at the interface to which these fragments were addressed which might not be necessarily the input interface for some of the fragments.
Reassembled Succeeded	The number of IPv6 datagrams successfully reassembled. Note that this counter is incremented at the interface to which these datagrams were addressed which might not be necessarily the input interface for some of the fragments.
Reassembled Failed	The number of failures detected by the IPv6 re-assembly algorithm (for whatever reason: timed out, errors, etc.). Note that this is not necessarily a count of discarded IPv6 fragments since some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by combining them as they are received. This counter is incremented at the interface to which these fragments were addressed which might not be necessarily the input interface for some of the fragments.

Show IPv6 Statistics - Display 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 ipv6IfStatsOutForwDatagrams.
Discards	The number of output IPv6 datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (e.g., for lack of buffer space). Note that this counter would include datagrams counted in ipv6IfStatsOutForwDatagrams if any such packets met this (discretionary) discard criterion.
No Routes	The number of input datagrams discarded because no route could be found to transmit them to their destination.
Generated Fragments	The number of output datagram fragments that have been generated as a result of fragmentation at this output interface.
Fragment Succeeded	The number of IPv6 datagrams that have been successfully fragmented at this output interface.
Fragment Failed	The number of IPv6 datagrams that have been discarded because they needed to be fragmented at this output interface but could not be.
ICMPv6 Statistics	
<i>ICMPv6 Received</i>	
Input	The total number of ICMP messages received by the interface which includes all those counted by ipv6IfIcmpInErrors. Note that this interface is the interface to which the ICMP messages were addressed which may not be necessarily the input interface for the messages.
Errors	The number of ICMP messages which the interface received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, etc.)
Destination Unreachable Messages	The number of ICMP Destination Unreachable messages received by the interface.
Packet Too Big Messages	The number of ICMP Packet Too Big messages received by the interface.
Time Exceeded Messages	The number of ICMP Time Exceeded messages received by the interface.
Parameter Problem Messages	The number of ICMP Parameter Problem messages received by the interface.
Echo Request Messages	The number of ICMP Echo (request) messages received by the interface.
Echo Reply Messages	The number of ICMP Echo Reply messages received by the interface.
Router Solicit Messages	The number of ICMP Router Solicit messages received by the interface.
Router Advertisement Messages	The number of ICMP Router Advertisement messages received by the interface.
Neighbor Solicit Messages	The number of ICMP Neighbor Solicit messages received by the interface.
Neighbor Advertisement Messages	The number of ICMP Neighbor Advertisement messages received by the interface
Redirect Messages	The number of Redirect messages received by the interface.
Group Membership Query Messages	The number of ICMPv6 Group Membership Query messages received by the interface.
Group Membership Response Messages	The number of ICMPv6 Group Membership Response messages received by the interface.
Group Membership Reduction Messages	The number of ICMPv6 Group Membership Reduction messages received by the interface.
Multicast Listener Discovery Version 2 Reports	The number of MLDv2 reports received by the interface.
<i>ICMPv6 Transmitted</i>	
Output	The total number of ICMP messages which this interface attempted to send. Note that this counter includes all those counted by icmpOutErrors.
Destination Unreachable Messages	The number of ICMP Destination Unreachable messages sent by the interface.
Packet Too Big Messages	The number of ICMP Packet Too Big messages sent by the interface.
Time Exceeded Messages	The number of ICMP Time Exceeded messages sent by the interface.
Echo Request Messages	The number of ICMP Echo (request) messages sent by the interface.
Echo Reply Messages	The number of ICMP Echo Reply messages sent by the interface.
Router Solicit Messages	The number of ICMP Router Solicit messages sent by the interface.
Router Advertisement Messages	The number of ICMP Router Advertisement messages sent by the interface.

Show IPv6 Statistics - Display Description	
Neighbor Solicit Messages	The number of ICMP Neighbor Solicit messages sent by the interface.
Neighbor Advertisement Messages	The number of ICMP Neighbor Advertisement messages sent by the interface
Redirect Messages	The number of Redirect messages sent by the interface.
Group Membership Query Messages	The number of ICMPv6 Group Membership Query messages sent by the interface.
Group Membership Response Messages	The number of ICMPv6 Group Membership Response messages sent by the interface.
Group Membership Reduction Messages	The number of ICMPv6 Group Membership Reduction messages sent by the interface.
Multicast Listener Discovery Version 2 Reports	The number of MLDv2 reports sent by the interface.
UDP Statistics	
Input	The total number of UDP datagrams delivered to UDP users.
No Port Errors	The total number of received UDP datagrams for which there was no application at the destination port.
Other Errors	The number of received UDP datagrams that could not be delivered for reasons other than the lack of an application at the destination port.
Output	The total number of UDP datagrams sent from this entity.

Perform these steps to show the IPv6 statistics:

1. Click **System > IPv6 Configuration**.
2. Select **Show Statistics** from the Action list.
3. Click **IPv6, ICMPv6, or UDP**.

The screenshot shows the 'System > IPv6 Configuration' page. The 'Action' dropdown is set to 'Show Statistics'. Under the 'Type' section, 'IPv6' is selected. The 'IPv6 Statistics' table displays the following data:

Total Received	0	Received Reassembly Succeeded	0
Received Header Errors	0	Received Reassembly Failed	0
Received Too Big Errors	0	Transmitted Forwards Datagrams	0
Received No Routes	0	Transmitted Requests	0
Received Address Errors	0	Transmitted Discards	0
Received Unknown Protocols	0	Transmitted No Routes	0
Received Truncated Packets	0	Transmitted Generated Fragments	0
Received Discards	0	Transmitted Fragment Succeeded	0
Received Delivers	0	Transmitted Fragment Failed	0
Received Reassembly Request Datagrams	0		

A 'Clear' button is located at the bottom of the statistics table.

FIG. 328 Showing IPv6 Statistics (IPv6)

The screenshot shows the 'System > IPv6 Configuration' page. The 'Action' dropdown is set to 'Show Statistics'. Under the 'Type' section, 'ICMPv6' is selected. The 'ICMPv6 Statistics' table displays the following data:

Received Input	0	Transmitted Output	0
Received Errors	0	Transmitted Destination Unreachable Messages	0
Received Destination Unreachable Messages	0	Transmitted Packet Too Big Messages	0
Received Packet Too Big Messages	0	Transmitted Time Exceeded Messages	0
Received Time Exceeded Messages	0	Transmitted Parameter Problem Message	0
Received Parameter Problem Messages	0	Transmitted Echo Request Messages	0
Received Echo Request Messages	0	Transmitted Echo Reply Messages	0
Received Echo Reply Messages	0	Transmitted Router Solicit Messages	0
Received Router Solicit Messages	0	Transmitted Router Advertisement Messages	0
Received Router Advertisement Messages	0	Transmitted Neighbor Solicit Messages	0
Received Neighbor Solicit Messages	0	Transmitted Neighbor Advertisement Messages	0
Received Neighbor Advertisement Messages	0	Transmitted Redirect Messages	0
Received Redirect Messages	0	Transmitted Group Membership Query Messages	0
Received Group Membership Query Messages	0	Transmitted Group Membership Response Messages	0
Received Group Membership Response Messages	0	Transmitted Group Membership Reduction Messages	0
Received Group Membership Reduction Messages	0		

A 'Clear' button is located at the bottom of the statistics table.

FIG. 329 Showing IPv6 Statistics (ICMPv6)

System > IPv6 Configuration

Action: Show Statistics

Type: IPv6 ICMPv6 UDP

UDP Statistics

Input	0
No Port Errors	0
Other Errors	0
Output	0

Clear

FIG. 330 Showing IPv6 Statistics (UDP)

Showing the MTU for Responding Destinations

Use the System > 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.

The following table lists the options on this page:

Show MTU - Display 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.

Perform these steps to show the MTU reported from other devices:

1. Click **System > IPv6 Configuration**.
2. Select **Show MTU** from the Action list.

System > IPv6 Configuration

Action: Show MTU

MTU Table Total: 2

MTU	Since	Destination Address
1400	00:04:21	5000:1::3
1280	00:04:50	FE80::203:A0FF:FED6:141D

FIG. 331 Showing Reported MTU Values

Appendix A: Software Specifications

Software Features

Management Authentication

Local, RADIUS, TACACS+, Port Authentication (802.1X), HTTPS, SSH, Port Security, IP Filter

General Security Measures

Access Control Lists (512 rules), Port Authentication (802.1X), MAC Authentication, Port Security, DHCP Snooping, IP Source Guard

Port Configuration

1000BASE-T: 10/100 Mbps at half/full duplex, 1000 Mbps at full duplex 1000BASE-SX/LX/ZX: 1000 Mbps at full duplex (SFP)

Flow Control

Full Duplex: IEEE 802.3-2005

Half Duplex: Back pressure

Storm Control

Broadcast, multicast, or unknown unicast traffic throttled above a critical threshold

Port Mirroring

10 sessions, one or more source ports to one destination port

Rate Limits

Input/Output Limits

Range configured per port

Port Trunking

Static trunks (Cisco EtherChannel compliant)

Dynamic trunks (Link Aggregation Control Protocol)

Spanning Tree Algorithm

Spanning Tree Protocol (STP, IEEE 802.1D-2004)

Rapid Spanning Tree Protocol (RSTP, IEEE 802.1D-2004)

Multiple Spanning Tree Protocol (MSTP, IEEE 802.1D-2004)

VLAN Support

Up to 4094 groups; port-based, protocol-based, tagged (802.1Q), voice VLANs, MAC-based

Class of Service

Supports four levels of priority

Strict, Weighted Round Robin (WRR), or a combination of strict and weighted queuing

Layer 3/4 priority mapping: IP DSCP

Quality of Service

DiffServ* supports class maps, policy maps, and service policies

* - Only supported for IPv4

Multicast Filtering

IGMP Snooping (Layer 2 IPv4)

MLD Snooping (Layer 2 IPv6)

IP Routing

ARP, CIDR (Classless Inter-Domain Routing)

Additional Features

BOOTP Client

DHCP Client, Option 82,

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 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

ARP (RFC 826)

DHCP Client (RFC 2131)

HTTPS

ICMP (RFC 792)

IGMP (RFC 1112)

IGMPv2 (RFC 2236)

IGMP Proxy (RFC 4541)

IPv4 IGMP (RFC 3228)

MLD Snooping (RFC 4541)

NTP (RFC 1305)

RADIUS+ (RFC 2618)

RMON (RFC 2819 groups 1,2,3,9)

SNMP (RFC 1157)

SNMPv2c (RFC 1901, 2571)

SNMPv3 (RFC DRAFT 2273, 2576, 3410, 3411, 3413, 3414, 3415)

SNTP (RFC 2030)

SSH (Version 2.0)

TELNET (RFC 854, 855, 856)

TFTP (RFC 1350)

Management Information Bases

Bridge MIB (RFC 1493)

Differentiated Services MIB (RFC 3289)

DNS Resolver MIB (RFC 1612)

Entity MIB (RFC 2737)

Ether-like MIB (RFC 2665)

Extended Bridge MIB (RFC 2674)

Extensible SNMP Agents MIB (RFC 2742)

Forwarding Table MIB (RFC 2096)
IGMP MIB (RFC 2933)
Interface Group MIB (RFC 2233)
Interfaces Evolution MIB (RFC 2863)
IP MIB (RFC 2011)
IP Forwarding Table MIB (RFC 2096)
IP Multicasting related MIBs
IPV6-MIB (RFC 2065)
IPV6-ICMP-MIB (RFC 2066)
IPV6-TCP-MIB (RFC 2052)
IPV6-UDP-MIB (RFC2054)
Link Aggregation MIB (IEEE 802.3ad)
MAU MIB (RFC 3636)
MIB II (RFC 1213)
NTP (RFC 1305)
P-Bridge MIB (RFC 2674P)
Port Access Entity MIB (IEEE 802.1X)
Port Access Entity Equipment MIB
Power Ethernet MIB (RFC 3621)
Private MIB
Q-Bridge MIB (RFC 2674Q)
Quality of Service MIB
RADIUS Accounting Server MIB (RFC 2621)
RADIUS Authentication Client MIB (RFC 2619)
RMON MIB (RFC 2819)
RMON II Probe Configuration Group (RFC 2021 , partial implementation)
SNMP Community MIB (RFC 3584)
SNMP Framework MIB (RFC 3411)
SNMP-MPD MIB (RFC 3412)
SNMP Target MIB, SNMP Notification MIB (RFC 3413)
SNMP User-Based SM MIB (RFC 3414)
SNMP View Based ACM MIB (RFC 3415)
SNMPv2 IP MIB (RFC 2011)
TACACS+ Authentication Client MIB
TCP MIB (RFC 2012)
Trap (RFC 1215)
UDP MIB (RFC 2013)

Appendix B: Troubleshooting

Diagnosing LED Indicators

LED Indicators	
LED Status	Action
Power LED is Off	<ul style="list-style-type: none"> • Check connections between the switch, the power cord, and the AC power outlet. • Check the AC power outlet is supplying 110-240VAC. • Contact your dealer for assistance.
DiagLED is blinking amber	<ul style="list-style-type: none"> • Power cycle the switch to try and clear the condition. • If the condition does not clear, contact your dealer for assistance.
Diag LED is blinking amber with PoE Mode button pressed	<ul style="list-style-type: none"> • Turn off or unplug PoE devices until the condition clears. • If the condition does not clear, contact your dealer for assistance.
Link LED is Off	<ul style="list-style-type: none"> • Verify that the switch and attached device are powered on. • Check the cable connectors are firmly plugged into both the switch and corresponding device. • If the switch is installed in a rack, check the connections to the punch-down block and patch panel. • Verify that the proper cable type is used and its length does not exceed specified limits. • Check the attached adapter and cable connections for possible defects. Replace the defective cable if necessary.

System Self-Diagnostic Test Failure

If the Diag LED indicates a failure of the system power-on-self-test (POST), you can use a console connection to view the POST results. The POST results may indicate a failed component or help troubleshoot the problem. For more information on connecting to the console port and using the CLI, refer to the CLI Reference Guide.

Note a POST failure normally indicates a serious hardware fault that cannot be rectified or worked around. If you encounter a POST failure, you should contact your dealer for assistance.

Power and Cooling Problems

If a power indicator does not turn on when the power cord is plugged in, you may have a problem with the power outlet, power cord, or internal power supply.

However, if the switch shuts down after operating for a continuous period, check for loose power connections, power losses or surges at the power outlet. If you still cannot isolate the problem, the internal power supply may be defective.

Installation

Verify that all system components have been properly installed. If one or more components appear to be malfunctioning (such as the power cord or network cabling), test them in an alternate environment where you are sure that all the other components are functioning properly.

In-Band Access

You can access the management agent in the switch through a connection to any port using Telnet, a web browser, or other network management software tools. However, you must first configure the switch with a valid IP address, subnet mask, and default gateway. If you have trouble establishing a link to the management agent, check to see if you have a valid network connection. Then verify that you entered the correct IP address. Also, be sure the switch port has not been disabled. If it has not been disabled, then check the network cabling that connects your remote location to the switch.

Problems Accessing the Management Interface

Troubleshooting Chart	
Symptom	Action
Cannot connect using Telnet, web browser, or SNMP software	<ul style="list-style-type: none"> • Be sure the switch is powered on. • Check network cabling between the management station and the switch. Make sure the ends are properly connected and there is no damage to the cable. Test the cable if necessary. • Check that you have a valid network connection to the switch and that the port you are using has not been disabled. • Be sure you have configured the VLAN interface through which the management station is connected with a valid IP address, subnet mask and default gateway. • Be sure the management station has an IP address in the same subnet as the switch's IP interface to which it is connected. • If you are trying to connect to the switch via the IP address for a tagged VLAN group, your management station, and the ports connecting intermediate switches in the network, must be configured with the appropriate tag. • If you cannot connect using Telnet, you may have exceeded the maximum number of concurrent Telnet/SSH sessions permitted. Try connecting again at a later time.
Cannot connect using Secure Shell	<ul style="list-style-type: none"> • If you cannot connect using SSH, you may have exceeded the maximum number of concurrent Telnet/SSH sessions permitted. Try connecting again at a later time. • Be sure the control parameters for the SSH server are properly configured on the switch, and that the SSH client software is properly configured on the management station. • Be sure you have generated both an RSA and DSA public key on the switch, exported this key to the SSH client, and enabled SSH service. Try using another SSH client or check for updates to your SSH client application. • Be sure you have set up an account on the switch for each SSH user, including user name, authentication level, and password. • Be sure you have imported the client's public key to the switch (if public key authentication is used).
Cannot access the on-board configuration program via a serial port connection	<ul style="list-style-type: none"> • Check to see if you have set the terminal emulator program to VT100 compatible, 8 data bits, 1 stop bit, no parity, and the baud rate set to 115200 bps. • Verify that you are using the DB-9 null-modem serial cable supplied with the switch. If you use any other cable, be sure that it conforms to the pin-out connections provided in the Installation Guide.
Forgot or lost the password	<ul style="list-style-type: none"> • Contact your local distributor.

Using System Logs

If a fault does occur, refer to the Installation Guide to ensure that the problem you encountered is actually caused by the switch. If the problem appears to be caused by the switch, follow these steps:

1. Enable logging.
2. Set the error messages reported to include all categories.
3. Enable SNMP.
4. Enable SNMP traps.
5. Designate the SNMP host that is to receive the error messages.
6. Repeat the sequence of commands or other actions that lead up to the error.
7. Make a list of the commands or circumstances that led to the fault. Also make a list of any error messages displayed.
8. Set up your terminal emulation software so that it can capture all console output to a file. Then enter the "show tech-support" command to record all system settings in this file.
9. Contact your distributor's service engineer, and send a detailed description of the problem, along with the file used to record your system settings.

For example:

```
Console(config)#logging on
Console(config)#logging history flash 7
Console(config)#snmp-server host 192.168.1.23
.
.
.
```



© 2016 Harman. All rights reserved. Metreau, NetLinx, AMX, AV FOR AN IT WORLD, HARMAN, and their respective logos are registered trademarks of HARMAN. Oracle, Java and any other company or brand name referenced may be trademarks/registered trademarks of their respective companies. AMX does not assume responsibility for errors or omissions. AMX also reserves the right to alter specifications without prior notice at any time. The AMX Warranty and Return Policy and related documents can be viewed/downloaded at www.amx.com.
3000 RESEARCH DRIVE, RICHARDSON, TX 75082 AMX.com | 800.222.0193 | 469.624.8000 | +1.469.624.7400 | fax 469.624.7153
AMX (UK) LTD, AMX by HARMAN - Unit C, Auster Road, Clifton Moor, York, YO30 4GD United Kingdom • +44 1904-343-100 • www.amx.com/eu/

Last Revised:
5/10/2016