



@server

Cisco Systems Intelligent Gigabit Ethernet Switch
Module for the IBM @server BladeCenter

Software Configuration Guide

Cisco IOS Release 12.1(14)AY

Note: Before using this information and the product it supports, read the general information in Appendix C. "Getting help and technical assistance" and Appendix D. "Notices".

Second Edition (June 2004)

© Copyright International Business Machines Corporation 2004. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Preface	13
Audience	13
Purpose	13
Conventions	13
Related Publications	14
Chapter 1. Overview	17
Features	17
Management Options	21
Management Interface Options	21
Advantages of Using CMS and Clustering Switches	21
Network Configuration Examples	22
Where to Go Next	23
Chapter 2. Using the Command-Line Interface	25
Cisco IOS Command Modes	25
Getting Help	27
Specifying Ports in Interface Configuration Mode	27
Abbreviating Commands	28
Using no and default Forms of Commands	28
Understanding CLI Messages	28
Using Command History	29
Changing the Command History Buffer Size	29
Recalling Commands	29
Disabling the Command History Feature	30
Using Editing Features	30
Enabling and Disabling Editing Features	30
Editing Commands through Keystrokes	31
Editing Command Lines that Wrap	32
Searching and Filtering Output of show and more Commands	33
Accessing the CLI	33
Accessing the CLI from a Browser	33
Chapter 3. Getting Started with CMS	35
Launching CMS	35
Features	36
Front Panel View	38
Cluster Tree	39
Front-Panel Images	40
Port Modes and LEDs	41
VLAN Membership Modes	41
Topology View	42
Topology Icons and Labels	43
Device and Link Labels	44
Colors in the Topology View	44
Topology Display Options	45
Menus and Toolbar	45
Menu Bar	45
Toolbar	50
Front Panel View Popup Menu	51
Topology View Popup Menu	52
Device Popup Menu	53
Interaction Modes	54

Guide Mode	54
Expert Mode	54
Wizards	54
Tool Tips	55
Online Help	55
CMS Window Components	56
Host Name List	57
Tabs, Lists, and Tables	57
Filter Editor	58
Buttons	58
Green Border around a Field or Cell	58
Red Border around a Field	58
Accessing CMS	58
Access Modes in CMS	59
HTTP Access to CMS	60
Saving Your Configuration	60
Restoring Your Configuration	61
CMS Preferences	61
Using Different Versions of CMS	61
Where to Go Next	61
Chapter 4. Assigning the Switch IP Address and Default Gateway	63
Understanding the Boot Process	63
Assigning Switch Information	63
Default Switch Information	64
Manually Assigning IP Information	64
Checking and Saving the Running Configuration	65
Modifying the Startup Configuration	67
Default Boot Configuration	67
Specifying the Filename to Read and Write the System Configuration	68
Booting Manually	68
Booting a Specific Software Image	69
Controlling Environment Variables	69
Scheduling a Reload of the Software Image	71
Configuring a Scheduled Reload	71
Displaying Scheduled Reload Information	72
Chapter 5. Clustering Switches	73
Understanding Switch Clusters	73
Command Switch Characteristics	74
Standby Command Switch Characteristics	74
Candidate Switch and Member Switch Characteristics	74
Planning a Switch Cluster	75
Automatic Discovery of Cluster Candidates and Members	75
HSRP and Standby Command Switches	79
IP Addresses	82
Host Names	82
Passwords	83
SNMP Community Strings	83
TACACS+ and RADIUS	83
Access Modes in CMS	83
Management VLAN	84
Availability of Switch-Specific Features in Switch Clusters	85
Creating a Switch Cluster	85
Enabling a Command Switch	85
Adding Member Switches	86

Creating a Cluster Standby Group	87
Verifying a Switch Cluster	89
Using the CLI to Manage Switch Clusters	90
Using SNMP to Manage Switch Clusters	90
Chapter 6. Administering the Switch	93
Managing the System Time and Date	93
Understanding the System Clock	93
Understanding Network Time Protocol	94
Configuring NTP	95
Configuring Time and Date Manually	102
Configuring a System Name and Prompt	105
Default System Name and Prompt Configuration	106
Configuring a System Name	106
Configuring a System Prompt	106
Understanding DNS	107
Creating a Banner	108
Default Banner Configuration	109
Configuring a Message-of-the-Day Login Banner	109
Configuring a Login Banner	110
Managing the MAC Address Table	110
Building the Address Table	111
MAC Addresses and VLANs	111
Default MAC Address Table Configuration	111
Changing the Address Aging Time	111
Removing Dynamic Address Entries	112
Configuring MAC Address Notification Traps	112
Adding and Removing Static Address Entries	114
Displaying Address Table Entries	115
Managing the ARP Table	116
Chapter 7. Configuring Switch-Based Authentication	117
Preventing Unauthorized Access to Your Switch	117
Protecting Access to Privileged EXEC Commands	117
Default Password and Privilege Level Configuration	118
Setting or Changing a Static Enable Password	118
Protecting Enable and Enable Secret Passwords with Encryption	119
Changing a Telnet Password for a Terminal Line	121
Configuring Username and Password Pairs	121
Configuring Multiple Privilege Levels	122
Controlling Switch Access with TACACS+	124
Understanding TACACS+	125
TACACS+ Operation	126
Configuring TACACS+	127
Displaying the TACACS+ Configuration	131
Controlling Switch Access with RADIUS	131
Understanding RADIUS	132
RADIUS Operation	133
Configuring RADIUS	133
Displaying the RADIUS Configuration	145
Configuring the Switch for Local Authentication and Authorization	145
Configuring the Switch for Secure Shell	146
Understanding SSH	147
Cryptographic Software Image Guidelines	147
Configuring SSH	147

Chapter 8. Configuring 802.1X Port-Based Authentication	149
Understanding 802.1X Port-Based Authentication	149
Device Roles	149
Authentication Initiation and Message Exchange	151
Ports in Authorized and Unauthorized States	152
Supported Topologies	152
Using 802.1X with Port Security	153
Using 802.1X with VLAN Assignment	154
Using 802.1X with Guest VLAN	155
Configuring 802.1X Authentication	155
Default 802.1X Configuration	156
802.1X Configuration Guidelines	156
Enabling 802.1X Authentication	157
Configuring the Switch-to-RADIUS-Server Communication	159
Enabling Periodic Re-Authentication	160
Manually Re-Authenticating a Client Connected to a Port	160
Changing the Quiet Period	161
Changing the Switch-to-Client Retransmission Time	161
Setting the Switch-to-Client Frame-Retransmission Number	162
Configuring the Host Mode	162
Configuring a Guest VLAN	163
Resetting the 802.1X Configuration to the Default Values	164
Displaying 802.1X Statistics and Status	164
Chapter 9. Configuring the Switch Interfaces	165
Understanding Interface Types	165
Access Ports	165
Trunk Ports	166
Port-Based VLANs	166
EtherChannel Port Groups	167
Connecting Interfaces	167
Using the Interface Command	168
Procedures for Configuring Interfaces	168
Configuring a Range of Interfaces	169
Configuring and Using Interface-Range Macros	171
Configuring Ethernet Interfaces	173
Default Ethernet Interface Configuration	173
Configuring Interface Speed and Duplex Mode	174
Configuring IEEE 802.3x Flow Control on Gigabit Ethernet Ports	176
Adding a Description for an Interface	177
Monitoring and Maintaining the Interfaces	178
Monitoring Interface and Controller Status	178
Clearing and Resetting Interfaces and Counters	180
Shutting Down and Restarting the Interface	181
Chapter 10. Configuring STP	183
Understanding Spanning-Tree Features	183
STP Overview	183
Spanning-Tree Topology and BPDUs	184
Bridge ID, Switch Priority, and Extended System ID	185
Spanning-Tree Interface States	186
How a Switch or Port Becomes the Root Switch or Root Port	188
Spanning Tree and Redundant Connectivity	189
Spanning-Tree Address Management	190
Accelerated Aging to Retain Connectivity	190
Spanning-Tree Modes and Protocols	190

Supported Spanning-Tree Instances	191
Spanning-Tree Interoperability and Backward Compatibility	191
STP and IEEE 802.1Q Trunks	191
Spanning Tree Considerations for Cisco Systems Intelligent Gigabit Ethernet Switch Modules	192
Configuring Spanning-Tree Features	192
Default Spanning-Tree Configuration	193
Spanning-Tree Configuration Guidelines	193
Changing the Spanning-Tree Mode	194
Disabling Spanning Tree	195
Configuring the Root Switch	196
Configuring a Secondary Root Switch	197
Configuring the Port Priority	198
Configuring the Path Cost	199
Configuring the Switch Priority of a VLAN	201
Configuring Spanning-Tree Timers	201
Displaying the Spanning-Tree Status	203
Chapter 11. Configuring MSTP	205
Understanding MSTP	205
Multiple Spanning-Tree Regions	206
IST, CIST, and CST	206
Hop Count	208
Boundary Ports	209
Interoperability with 802.1D STP	209
Understanding RSTP	210
Port Roles and the Active Topology	210
Rapid Convergence	211
Synchronization of Port Roles	212
Bridge Protocol Data Unit Format and Processing	213
Topology Changes	214
Configuring MSTP Features	215
Default MSTP Configuration	215
MSTP Configuration Guidelines	216
Specifying the MST Region Configuration and Enabling MSTP	217
Configuring the Root Switch	218
Configuring a Secondary Root Switch	219
Configuring the Port Priority	220
Configuring the Path Cost	221
Configuring the Switch Priority	222
Configuring the Hello Time	223
Configuring the Forwarding-Delay Time	224
Configuring the Maximum-Aging Time	224
Configuring the Maximum-Hop Count	224
Specifying the Link Type to Ensure Rapid Transitions	225
Restarting the Protocol Migration Process	225
Displaying the MST Configuration and Status	226
Chapter 12. Configuring Optional Spanning-Tree Features	227
Understanding Optional Spanning-Tree Features	227
Understanding Port Fast	227
Understanding BPDU Guard	228
Understanding BPDU Filtering	228
Understanding UplinkFast	229
Understanding BackboneFast	231
Understanding EtherChannel Guard	233

Understanding Root Guard	233
Understanding Loop Guard	234
Configuring Optional Spanning-Tree Features	234
Default Optional Spanning-Tree Configuration	235
Optional Spanning-Tree Configuration Guidelines	235
Enabling Port Fast	235
Enabling BPDU Guard	236
Enabling BPDU Filtering	237
Enabling UplinkFast for Use with Redundant Links	238
Enabling BackboneFast	239
Enabling EtherChannel Guard	239
Enabling Root Guard	240
Enabling Loop Guard	240
Displaying the Spanning-Tree Status	241
Chapter 13. Configuring VLANs	243
Understanding VLANs	243
Supported VLANs	244
VLAN Port Membership Modes	244
Configuring Normal-Range VLANs	245
Token Ring VLANs	246
Normal-Range VLAN Configuration Guidelines	247
VLAN Configuration Mode Options	247
Saving VLAN Configuration	248
Default Ethernet VLAN Configuration	249
Creating or Modifying an Ethernet VLAN	249
Deleting a VLAN	251
Assigning Static-Access Ports to a VLAN	251
Configuring Extended-Range VLANs	252
Default VLAN Configuration	253
Extended-Range VLAN Configuration Guidelines	253
Creating an Extended-Range VLAN	253
Displaying VLANs	254
Configuring VLAN Trunks	255
Trunking Overview	255
Default Layer 2 Ethernet Interface VLAN Configuration	257
Configuring an Ethernet Interface as a Trunk Port	258
Load Sharing Using STP	262
Configuring VMPS	266
Understanding VMPS	266
Default VMPS Configuration	269
VMPS Configuration Guidelines	269
Configuring the VMPS Client	270
Monitoring the VMPS	272
Troubleshooting Dynamic Port VLAN Membership	273
VMPS Configuration Example	273
Chapter 14. Configuring VTP	275
Understanding VTP	275
The VTP Domain	275
VTP Modes	276
VTP Advertisements	277
VTP Version 2	278
VTP Pruning	278
Configuring VTP	280
Default VTP Configuration	280

VTP Configuration Options	280
VTP Configuration Guidelines	281
Configuring a VTP Server	282
Configuring a VTP Client	284
Disabling VTP (VTP Transparent Mode)	285
Enabling VTP Version 2	286
Enabling VTP Pruning	286
Adding a VTP Client Switch to a VTP Domain	287
Monitoring VTP	288
Chapter 15. Configuring IGMP Snooping and MVR	289
Understanding IGMP Snooping	289
Joining a Multicast Group	290
Leaving a Multicast Group	292
Immediate-Leave Processing	292
Source-Only Networks	292
Configuring IGMP Snooping	293
Default IGMP Snooping Configuration	293
Enabling or Disabling IGMP Snooping	293
Setting the Snooping Method	294
Configuring a Multicast Router Port	295
Configuring a Host Statically to Join a Group	296
Enabling IGMP Immediate-Leave Processing	297
Disabling IP Multicast-Source-Only Learning	297
Configuring the Aging Time	298
Displaying IGMP Snooping Information	299
Understanding Multicast VLAN Registration	301
Using MVR in a Multicast Television Application	302
Configuring MVR	303
Default MVR Configuration	303
MVR Configuration Guidelines and Limitations	304
Configuring MVR Global Parameters	304
Configuring MVR Interfaces	305
Displaying MVR Information	307
Configuring IGMP Filtering	308
Default IGMP Filtering Configuration	309
Configuring IGMP Profiles	309
Applying IGMP Profiles	310
Setting the Maximum Number of IGMP Groups	311
Displaying IGMP Filtering Configuration	312
Chapter 16. Configuring Port-Based Traffic Control	315
Configuring Storm Control	315
Understanding Storm Control	315
Default Storm Control Configuration	316
Enabling Storm Control	316
Disabling Storm Control	316
Configuring Protected Ports	317
Configuring Port Security	318
Understanding Port Security	318
Default Port Security Configuration	320
Port Security Configuration Guidelines	320
Enabling and Configuring Port Security	321
Enabling and Configuring Port Security Aging	323
Displaying Port-Based Traffic Control Settings	325

Chapter 17. Configuring UDLD	327
Understanding UDLD	327
Configuring UDLD	328
Default UDLD Configuration	328
Enabling UDLD Globally	329
Enabling UDLD on an Interface	329
Resetting an Interface Shut Down by UDLD	330
Displaying UDLD Status	330
Chapter 18. Configuring CDP	331
Understanding CDP	331
Configuring CDP	331
Default CDP Configuration	331
Configuring the CDP Characteristics	332
Disabling and Enabling CDP	333
Disabling and Enabling CDP on an Interface	333
Monitoring and Maintaining CDP	334
Chapter 19. Configuring SPAN and RSPAN	337
Understanding SPAN and RSPAN	337
SPAN and RSPAN Concepts and Terminology	338
SPAN and RSPAN Interaction with Other Features	341
SPAN and RSPAN Session Limits	342
Default SPAN and RSPAN Configuration	342
Configuring SPAN	342
SPAN Configuration Guidelines	342
Creating a SPAN Session and Specifying Ports to Monitor	343
Creating a SPAN Session and Enabling Ingress Traffic	344
Removing Ports from a SPAN Session	346
Configuring RSPAN	347
RSPAN Configuration Guidelines	347
Creating an RSPAN Session	348
Creating an RSPAN Destination Session	349
Removing Ports from an RSPAN Session	350
Displaying SPAN and RSPAN Status	351
Chapter 20. Configuring RMON	353
Understanding RMON	353
Configuring RMON	354
Default RMON Configuration	354
Configuring RMON Alarms and Events	354
Configuring RMON Collection on an Interface	356
Displaying RMON Status	357
Chapter 21. Configuring System Message Logging	359
Understanding System Message Logging	359
Configuring System Message Logging	359
System Log Message Format	360
Default System Message Logging Configuration	361
Disabling and Enabling Message Logging	361
Setting the Message Display Destination Device	362
Synchronizing Log Messages	363
Enabling and Disabling Timestamps on Log Messages	364
Enabling and Disabling Sequence Numbers in Log Messages	365
Defining the Message Severity Level	365
Limiting Syslog Messages Sent to the History Table and to SNMP	366

Configuring UNIX Syslog Servers	367
Displaying the Logging Configuration	369
Chapter 22. Configuring Network Security with ACLs	371
Understanding ACLs	371
Handling Fragmented and Unfragmented Traffic	372
Understanding Access Control Parameters	373
Guidelines for Applying ACLs to Physical Interfaces	375
Configuring ACLs	375
Unsupported Features	376
Creating Standard and Extended IP ACLs	376
Creating Named MAC Extended ACLs	386
Creating MAC Access Groups	387
Applying ACLs to Terminal Lines or Physical Interfaces	387
Applying ACLs to a Physical Interface	388
Displaying ACL Information	388
Displaying ACLs	389
Displaying Access Groups	389
Examples for Compiling ACLs	390
Numbered ACL Examples	391
Extended ACL Examples	391
Named ACL Example	392
Commented IP ACL Entry Examples	392
Chapter 23. Configuring SNMP	395
Understanding SNMP	395
SNMP Versions	395
SNMP Manager Functions	397
SNMP Agent Functions	397
SNMP Community Strings	397
Using SNMP to Access MIB Variables	398
SNMP Notifications	398
Configuring SNMP	399
Default SNMP Configuration	399
SNMP Configuration Guidelines	399
Disabling the SNMP Agent	400
Configuring Community Strings	400
Configuring SNMP Groups and Users	402
Configuring SNMP Notifications	404
Setting the Agent Contact and Location Information	407
Limiting TFTP Servers Used Through SNMP	407
SNMP Examples	408
Displaying SNMP Status	409
Chapter 24. Configuring QoS	411
Understanding QoS	411
Basic QoS Model	413
Classification	413
Policing and Marking	415
Mapping Tables	416
Queueing and Scheduling	416
Configuring Auto-QoS	417
Generated Auto-QoS Configuration	418
Effects of Auto-QoS on the Configuration	420
Configuration Guidelines	420
Enabling Auto-QoS for VoIP	420

Displaying Auto-QoS Information	421
Configuring Standard QoS	421
Default Standard QoS Configuration	422
Configuration Guidelines	422
Configuring Classification Using Port Trust States	423
Configuring a QoS Policy	428
Configuring CoS Maps	436
Configuring the Egress Queues	438
Displaying Standard QoS Information	440
Standard QoS Configuration Examples	440
QoS Configuration for the Existing Wiring Closet	441
QoS Configuration for the Intelligent Wiring Closet	442
Chapter 25. Configuring EtherChannels	443
Understanding EtherChannels	443
Understanding Port-Channel Interfaces	444
Understanding the Port Aggregation Protocol and Link Aggregation Protocol	444
Understanding Load Balancing and Forwarding Methods	446
Configuring EtherChannels	448
Default EtherChannel Configuration	449
EtherChannel Configuration Guidelines	449
Configuring Layer 2 EtherChannels	450
Configuring EtherChannel Load Balancing	452
Configuring the PAgP Learn Method and Priority	453
Configuring the LACP Port Priority	453
Configuring Hot Standby Ports	454
Configuring the LACP System Priority	454
Displaying EtherChannel, PAgP, and LACP Status	455
Chapter 26. Troubleshooting	457
Preventing Autonegotiation Mismatches	457
Diagnosing Connectivity Problems	457
Using Ping	457
Using Layer 2 Traceroute	459
Using Debug Commands	460
Enabling Debugging on a Specific Feature	461
Enabling All-System Diagnostics	461
Redirecting Debug and Error Message Output	462
Using the debug autoqos Command	462
Using the crashinfo File	463
Appendix A. Supported MIBs	465
MIB List	465
Accessing the MIB Files	466
Appendix B. Working with the Cisco IOS File System, Configuration Files, and Software Images	467
Working with the Flash File System	467
Displaying Available File Systems	467
Setting the Default File System	468
Displaying Information about Files on a File System	469
Changing Directories and Displaying the Working Directory	469
Creating and Removing Directories	469
Copying Files	470
Deleting Files	471
Creating, Displaying, and Extracting tar Files	471

Displaying the Contents of a File	473
Working with Configuration Files	474
Guidelines for Creating and Using Configuration Files	474
Configuration File Types and Location	475
Creating a Configuration File By Using a Text Editor	475
Copying Configuration Files By Using TFTP	475
Copying Configuration Files By Using FTP	477
Copying Configuration Files By Using RCP	480
Clearing Configuration Information	484
Working with Software Images	484
Image Location on the Switch	485
tar File Format of Images on a Server or Cisco.com	485
Copying Image Files By Using TFTP	486
Copying Image Files By Using FTP	489
Copying Image Files By Using RCP	493
Appendix C. Getting help and technical assistance	499
Before you call	499
Using the documentation	499
Getting help and information from the World Wide Web	499
Software service and support	500
Hardware service and support	500
Appendix D. Notices	501
Edition notice	501
Trademarks	501

Preface

Audience

This guide is for the network manager responsible for configuring the Cisco Systems Intelligent Gigabit Ethernet Switch Module, hereafter referred to as the *switch*. Before using this guide, you should be familiar with the concepts and terminology of Ethernet and local area networking.

Purpose

This guide provides information about configuring and troubleshooting a switch or switch clusters. It includes descriptions of the management interface options and the features supported by the switch software.

Use this guide with other documents for information about these topics:

- Requirements—This guide assumes that you have met the hardware and software requirements and cluster compatibility requirements described in the release notes.
- Start-up information—This guide assumes that you have assigned switch IP information and passwords by using the browser setup program described in the *Cisco Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter Installation Guide*.
- Cluster Management Suite (CMS) information—This guide provides an overview of the CMS web-based, switch management interface. For information about CMS requirements and the procedures for browser and plug-in configuration and accessing CMS, refer to the release notes. For CMS field-level window descriptions and procedures, refer to the CMS online help.
- Cluster configuration—This guide provides information about planning for, creating, and maintaining switch clusters. Because configuring switch clusters is most easily performed through CMS, this guide does not provide the command-line interface (CLI) procedures. For the cluster commands, refer to the command reference for this release.
- CLI command information—This guide provides an overview for using the CLI. For complete syntax and usage information about the commands that have been specifically created or changed for the switches, refer to the command reference for this release.

Note: This guide does not repeat the concepts and CLI procedures provided in the standard Cisco IOS Release 12.1 documentation. For information about the standard Cisco IOS Release 12.1 commands, refer to the Cisco IOS documentation set available from the Cisco.com home page at **Service and Support > Technical Documents**. On the Cisco Product Documentation home page, select Release 12.1 from the Cisco IOS Software drop-down list.

- This guide does not describe system messages you might encounter. For more information, refer to the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Message Guide* for this release.

Conventions

This guide uses these conventions to convey instructions and information:

Command descriptions use these conventions:

- Commands and keywords are in **boldface** text.
- Arguments for which you supply values are in *italic*.
- Square brackets ([]) indicate optional elements.
- Braces ({ }) group required choices, and vertical bars (|) separate the alternative elements.
- Braces and vertical bars within square brackets ({ | }) indicate a required choice within an optional element.

Interactive examples use these conventions:

- Terminal sessions and system displays are in `screen` font.
- Information you enter is in **boldface screen** font.
- Nonprinting characters, such as passwords or tabs, are in angle brackets (< >).

Notes, cautions, and tips use these conventions and symbols:

Note: Means *reader take note*. Notes contain helpful suggestions or references to materials not contained in this manual.

Caution: Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Important: Means *the following will help you solve a problem*. The tips information might not be troubleshooting or even an action, but could be useful information.

Related Publications

In addition to this document, the following related documentation comes with the Gigabit Ethernet switch module:

- *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Release Notes*

Note: Switch requirements and procedures for initial configurations and software upgrades tend to change and therefore appear only in the release notes. Before installing, configuring, or upgrading the switch, refer to the release notes for the latest information.

- *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference* (order number TBD)

This document is in PDF form on the IBM *BladeCenter Documentation* CD. It includes:

- Command line interface (CLI) modes
- Command line interface commands and examples
- Syntax description
- Defaults
- Command history
- Usage guidelines
- Related commands

- *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Message Guide*

This document is in PDF on the IBM *BladeCenter Documentation* CD. It contains information about the switch-specific system messages. During operation, the system software sends these messages to the console or logging server on another system. Not all system messages indicate problems with the system.

Some messages are informational, while others can help diagnose problems with communication lines, internal hardware, or the system software. This document also includes error messages that display when the system fails.

- *Cisco Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter Installation Guide*

This document contains installation and configuration instructions for the Gigabit Ethernet switch module. This document also provides general information about your Gigabit Ethernet switch module, including warranty information, and how to get help. This document is also on the IBM *BladeCenter Documentation CD*.

- *eServer BladeCenter Type 8677 Installation and User's Guide*

This document is in PDF on the IBM *BladeCenter Documentation CD*. It contains general information about your BladeCenter unit, including:

- Information about features
- How to set up, cable, and start the BladeCenter unit
- How to install options in the BladeCenter unit
- How to configure the BladeCenter unit
- How to perform basic troubleshooting of the BladeCenter unit
- How to get help

- *BladeCenter Management Module User's Guide*

This document is in PDF on the IBM *BladeCenter Documentation CD*. It provides general information about the management module, including:

- Information about features
- How to start the management module
- How to install the management module
- How to configure and use the management module

- *BladeCenter HS20 Installation and User's Guide (for each blade server type)*

These documents are in PDF on the IBM *BladeCenter Documentation CD*. Each provides general information about a blade server, including:

- Information about features
- How to set up and start your blade server
- How to install options in your blade server
- How to configure your blade server
- How to install an operating system on your blade server
- How to perform basic troubleshooting of your blade server
- How to get help

- Cisco IOS Release 12.1 documentation at <http://www.cisco.com/en/US/products/sw/iosswrel/ps1831/index.html>
- Cisco IOS Release 12.2 documentation at <http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/index.html>

Chapter 1. Overview

This chapter provides these topics about the switch software:

- Features, on page 17
- Management Options, on page 21
- Network Configuration Examples, on page 22
- Where to Go Next, on page 23

Features

This section describes the features supported in this release:

Ease of Use and Ease of Deployment

- Cluster Management Suite (CMS) software for simplifying switch and switch cluster management through a web browser, such as Netscape Communicator or Microsoft Internet Explorer, from anywhere in your intranet
- Switch clustering technology used with CMS for
 - Unified configuration, monitoring, authentication, and software upgrade of multiple switches (refer to the release notes for a list of eligible cluster members).
 - Automatic discovery of candidate switches and creation of clusters of up to 16 switches that can be managed through a single IP address.
 - Extended discovery of cluster candidates that are not directly connected to the command switch.
- Hot Standby Router Protocol (HSRP) for command-switch redundancy. The redundant command switches used for HSRP must have compatible software releases.

Note: See the “Advantages of Using CMS and Clustering Switches” section on page 21. For the CMS, software, and browser requirements and for the cluster hardware and software requirements, refer to the release notes.

Performance

- Autosensing of speed on the 10/100/1000 Mbps ports and autonegotiation of duplex mode on the external ports for optimizing bandwidth
- IEEE 802.3x flow control on Gigabit Ethernet ports operating in full-duplex mode
- Fast EtherChannel and Gigabit EtherChannel for enhanced fault tolerance and for providing up to 4 Gbps of bandwidth between switches, routers, and servers
- Support for frame sizes from 64 to 1530 bytes.
- Per-port broadcast storm control for preventing faulty end stations from degrading overall system performance with broadcast storms
- Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP) for automatic creation of EtherChannel links
- Internet Group Management Protocol (IGMP) snooping support to limit flooding of IP multicast traffic
- Multicast VLAN registration (MVR) to continuously send multicast streams in a multicast VLAN while isolating the streams from subscriber VLANs for bandwidth and security reasons
- IGMP filtering for controlling the set of multicast groups to which hosts on a switch port can belong

- Protected port (private VLAN edge port) option for restricting the forwarding of traffic to designated ports on the same switch
- Dynamic address learning for enhanced security

Manageability

- Address Resolution Protocol (ARP) for identifying a switch through its IP address and its corresponding MAC address
- Cisco Discovery Protocol (CDP) versions 1 and 2 for network topology discovery and mapping between the switch and other Cisco devices on the network
- Network Time Protocol (NTP) for providing a consistent timestamp to all switches from an external source
- Directed unicast requests to a Trivial File Transfer Protocol (TFTP) server for obtaining software upgrades from a TFTP server
- Default configuration storage in flash memory to ensure that the switch can be connected to a network and can forward traffic with minimal user intervention
- In-band management access through a CMS web-based session
- In-band management access through up to 16 simultaneous Telnet connections for multiple command-line interface (CLI)-based sessions over the network
- In-band management access through up to five simultaneous, encrypted Secure Shell (SSH) connections for multiple CLI-based sessions over the network (only available in the cryptographic software image)
- In-band management access through SNMP versions 1, 2c, and 3 get and set requests
- Out-of-band management access to a remote terminal through a serial connection and a modem

Note: For additional descriptions of the management interfaces, see the “Management Options” section on page 21.

Redundancy

- HSRP for command-switch redundancy
- UniDirectional link detection (UDLD) on all Ethernet ports for detecting and disabling unidirectional links caused by port faults
- IEEE 802.1D Spanning Tree Protocol (STP) for redundant backbone connections and loop-free networks. STP has these features:
 - Per-VLAN spanning-tree plus (PVST+) for balancing load across VLANs
 - Rapid PVST+ for balancing load across VLANs
 - UplinkFast and BackboneFast for fast convergence after a spanning-tree topology change and for achieving load balancing between redundant uplinks, including Gigabit uplinks
- IEEE 802.1s Multiple STP (MSTP) for grouping VLANs into a spanning-tree instance, and providing for multiple forwarding paths for data traffic and load balancing
- IEEE 802.1w Rapid STP (RSTP) for rapid convergence of the spanning tree by immediately transitioning root and designated ports to the forwarding state
- Optional spanning-tree features available in the PVST+, rapid PVST+, and MSTP modes:
 - Port Fast for eliminating the forwarding delay by enabling a port to immediately transition from the blocking state to the forwarding state
 - BPDU guard for shutting down Port Fast-enabled ports that receive BPDUs

- BPDU filtering for preventing a Port Fast-enabled port from sending or receiving BPDUs
- Root guard for preventing switches outside the network core from becoming the spanning-tree root
- Loop guard for preventing alternate or root ports from becoming designated ports because of a failure that leads to a unidirectional link

Note: The switch supports up to 64 spanning-tree instances.

VLAN Support

- The switches support 250 port-based VLANs for assigning users to VLANs associated with appropriate network resources, traffic patterns, and bandwidth
- The switch supports up to 4094 VLAN IDs to allow service provider networks to support the number of VLANs allowed by the IEEE 802.1Q standard
- IEEE 802.1Q trunking protocol on all ports for network moves, adds, and changes; management and control of broadcast and multicast traffic; and network security by establishing VLAN groups for high-security users and network resources
- VLAN Membership Policy Server (VMPS) for dynamic VLAN membership
- VLAN Trunking Protocol (VTP) pruning for reducing network traffic by restricting flooded traffic to links destined for stations receiving the traffic
- Dynamic Trunking Protocol (DTP) for negotiating trunking on a link between two devices and for negotiating the type of trunking encapsulation (802.1Q) to be used
- Voice VLAN for creating subnets for voice traffic from Cisco IP Phones
- VLAN 1 minimization to reduce the risk of spanning-tree loops or storms by allowing VLAN 1 to be disabled on any individual VLAN trunk link. With this feature enabled, no user traffic is sent or received. The switch CPU continues to send and receive control protocol frames.

Security

- Bridge protocol data unit (BPDU) guard for shutting down a Port Fast-configured port when an invalid configuration occurs
- Protected port option for restricting the forwarding of traffic to designated ports on the same switch
- Password-protected access (read-only and read-write access) to management interfaces (CMS and CLI) for protection against unauthorized configuration changes
- Port security option for limiting and identifying MAC addresses of the stations allowed to access the port
- Port security aging to set the aging time for secure addresses on a port
- Multilevel security for a choice of security level, notification, and resulting actions
- MAC-based port-level security for restricting the use of a switch port to a specific group of source addresses and preventing switch access from unauthorized stations
- Terminal Access Controller Access Control System Plus (TACACS+), a proprietary feature for managing network security through a TACACS server
- IEEE 802.1X port-based authentication to prevent unauthorized devices from gaining access to the network
- IEEE 802.1X port-based authentication with VLAN assignment for restricting 802.1X-authenticated users to a specified VLAN

- IEEE 802.1X port-based authentication with port security for authenticating the port and managing network access for all MAC addresses, including that of the client
- IEEE 802.1X port-based authentication with port security for controlling access to 802.1X multiple-host ports
- IEEE 802.1X port-based authentication with voice VLAN to permit an IP phone access to the voice VLAN irrespective of the authorized or unauthorized state of the port
- IEEE 802.1X port-based authentication with guest VLAN to provide limited services to non-802.1X-compliant users

Quality of Service and Class of Service

- Classification
 - IEEE 802.1p class of service (CoS) with eight priority queues on the Gigabit ports for prioritizing mission-critical and time-sensitive traffic from data, voice, and telephony applications
 - Support for IEEE 802.1p CoS scheduling for classification and preferential treatment of high-priority voice traffic
 - Trusted boundary (detect the presence of a Cisco IP Phone, trust the CoS value received, and ensure port security. If the IP phone is not detected, disable the trusted setting on the port and prevent misuse of a high-priority queue.)
- Policing
 - Traffic-policing policies on the switch port for allocating the amount of the port bandwidth to a specific traffic flow
 - Policing traffic flows to restrict specific applications or traffic flows to metered, predefined rates
 - Up to 60 policers on ingress Gigabit-capable Ethernet ports
Granularity of 8 Mbps on 10/100/1000 Mbps ports
 - Out-of-profile markdown for packets that exceed bandwidth utilization limits
- Egress Policing and Scheduling of Egress Queues—Four egress queues on all switch ports. Support for strict priority and weighted round-robin (WRR) CoS policies

Monitoring

- Switch LEDs that provide visual external port and switch status
- Switched Port Analyzer (SPAN) for traffic monitoring on any port or VLAN
- SPAN support of Intrusion Detection Systems (IDSs) to monitor, repel, and report network security violations
- Four groups (history, statistics, alarms, and events) of embedded remote monitoring (RMON) agents for network monitoring and traffic analysis
- MAC address notification for tracking the MAC addresses that the switch has learned or removed
- Syslog facility for logging system messages about authentication or authorization errors, resource issues, and time-out events
- Layer 2 traceroute to identify the physical path that a packet takes from a source device to a destination device

Management Options

The switches are designed for plug-and-play operation: you only need to assign basic IP information to the switch and connect it to the other devices in your network. If you have specific network needs, you can configure and monitor the switch—on an individual basis or as part of a switch cluster—through its various management interfaces.

This section discusses these topics:

- Management Interface Options, on page 21
- Advantages of Using CMS and Clustering Switches, on page 21

Management Interface Options

You can configure and monitor individual switches and switch clusters by using these interfaces:

- CMS—CMS is a graphical user interface that can be launched from anywhere in your network through a web browser such as Netscape Communicator or Microsoft Internet Explorer. CMS is already installed on the switch. Using CMS, you can configure and monitor a standalone switch, a specific cluster member, or an entire switch cluster. You can also display network topologies to gather link information and display switch images to modify switch and port level settings.

For more information about CMS, see Chapter 3 “Getting Started with CMS.”

- CLI—The switch Cisco IOS CLI software is enhanced to support desktop-switching features. You can configure and monitor the switch and switch cluster members from the CLI. You can access the CLI by using Telnet or SSH from a remote management station. You can also access the CLI through the serial port connector on the rear of the Cisco Intelligent Gigabit Ethernet Switch Module. See the *Cisco Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter Installation Guide*.

For more information about the CLI, see Chapter 2 “Using the Command-Line Interface.”

- SNMP—SNMP provides a means to monitor and control the switch and switch cluster members. You can manage switch configuration settings, performance, and security and collect statistics by using SNMP management applications such as CiscoWorks2000 LAN Management Suite (LMS) and HP OpenView.

You can manage the switch from an SNMP-compatible management station that is running platforms such as HP OpenView or SunNet Manager. The switch supports a comprehensive set of MIB extensions and four RMON groups.

For more information about using SNMP, see Chapter 23 “Configuring SNMP.”

Advantages of Using CMS and Clustering Switches

Using CMS and switch clusters can simplify and minimize your configuration and monitoring tasks. You can use Cisco switch clustering technology to manage up to 16 interconnected and supported switches through one IP address as if they were a single entity. This can conserve IP addresses if you have a limited number of them. CMS is the easiest interface to use and makes switch and switch cluster management accessible to authorized users from any PC on your network.

By using switch clusters and CMS, you can:

- Manage and monitor interconnected switches (refer to the release notes for a list of supported switches), regardless of their geographic proximity and

interconnection media, including Ethernet, Gigabit Ethernet, and Gigabit EtherChannel connections.


- Accomplish multiple configuration tasks from a single CMS window without needing to remember CLI commands to accomplish specific tasks.
- Apply actions from CMS to multiple ports and multiple switches at the same time to avoid re-entering the same commands for each individual port or switch. Here are some examples of globally setting and managing multiple ports and switches:
 - Port configuration such as speed and duplex settings
 - Port security settings
 - NTP, STP, VLAN, and quality of service (QoS) configurations
 - Inventory and statistic reporting and link and switch-level monitoring and troubleshooting
 - Group software upgrades
- View a topology of interconnected devices to identify existing switch clusters and eligible switches that can join a cluster. You can also use the topology to quickly identify link information between switches.
- Monitor real-time status of a switch or multiple switches from the LEDs on the front-panel images. The system, redundant power system (RPS), and port LED colors on the images are similar to those on the physical LEDs.
- Use an interactive mode that takes you step-by-step through configuring complex features such as VLANs, ACLs, and QoS.
- Use a wizard that prompts you to provide the minimum required information to configure complex features such as QoS priorities for video traffic, priority levels for data applications, and security.

For more information about CMS, see Chapter 3 “Getting Started with CMS.” For more information about switch clusters, see Chapter 5 “Clustering Switches.”

Network Configuration Examples

The following figures show three different network configurations.

Figure 1. Basic Configuration



15929

Figure 2. Trunking Configuration





Figure 3. Redundancy Configuration



Where to Go Next

Before configuring the switch, review these sections for start up information:

- Chapter 2 “Using the Command-Line Interface”
- Chapter 3 “Getting Started with CMS”
- Chapter 4 “Assigning the Switch IP Address and Default Gateway”

Chapter 2. Using the Command-Line Interface

This chapter describes the Cisco IOS command-line interface (CLI) that you can use to configure your switch. It contains these sections:

- Cisco IOS Command Modes, on page 25
- Getting Help, on page 27
- Specifying Ports in Interface Configuration Mode, on page 27
- Abbreviating Commands, on page 28
- Using no and default Forms of Commands, on page 28
- Understanding CLI Messages, on page 28
- Using Command History, on page 29
- Using Editing Features, on page 30
- Searching and Filtering Output of show and more Commands, on page 33
- Accessing the CLI, on page 33
- Accessing the CLI from a Browser, on page 33

Cisco IOS Command Modes

The Cisco IOS user interface is divided into many different modes. The commands available to you depend on which mode you are currently in. Enter a question mark (?) at the system prompt to obtain a list of commands available for each command mode.

When you start a session on the switch, you begin in user mode, often called user EXEC mode. Only a limited subset of the commands are available in user EXEC mode. For example, most of the user EXEC commands are one-time commands, such as **show** commands, which show the current configuration status, and **clear** commands, which clear counters or interfaces. The user EXEC commands are not saved when the switch reboots.

To have access to all commands, you must enter privileged EXEC mode. Normally, you must enter a password to enter privileged EXEC mode. From this mode, you can enter any privileged EXEC command or enter global configuration mode.

Using the configuration modes (global, interface, and line), you can make changes to the running configuration. If you save the configuration, these commands are stored and used when the switch reboots. To access the various configuration modes, you must start at global configuration mode. From global configuration mode, you can enter interface configuration mode and line configuration mode.

Table 1 describes the main command modes, how to access each one, the prompt you see in that mode, and how to exit the mode. The examples in the table use the host name *Switch*.

Table 1. Command Mode Summary

Mode	Access Method	Prompt	Exit Method	About This Mode
User EXEC	Begin a session with your switch.	Switch>	Enter logout or quit .	Use this mode to <ul style="list-style-type: none"> • Change terminal settings. • Perform basic tests. • Display system information.
Privileged EXEC	While in user EXEC mode, enter the enable command.	Switch#	Enter disable to exit.	Use this mode to verify commands that you have entered. Use a password to protect access to this mode.
Global configuration	While in privileged EXEC mode, enter the configure command.	Switch(config)#	To exit to privileged EXEC mode, enter exit or end , or press Ctrl-Z .	Use this mode to configure parameters that apply to the entire switch.
Config-vlan	While in global configuration mode, enter the vlan vlan-id command.	Switch(config-vlan)#	To exit to global configuration mode, enter the exit command. To return to privileged EXEC mode, press Ctrl-Z or enter end .	Use this mode to configure VLAN parameters. When VTP mode is transparent, you can create extended-range VLANs (VLAN IDs greater than 1005) and save configurations in the switch startup configuration file.
VLAN configuration	While in privileged EXEC mode, enter the vlan database command.	Switch(vlan)#	To exit to privileged EXEC mode, enter exit .	Use this mode to configure VLAN parameters for VLANs 1 to 1005 in the VLAN database.
Interface configuration	While in global configuration mode, enter the interface command (with a specific interface).	Switch(config-if)#	To exit to global configuration mode, enter exit . To return to privileged EXEC mode, press Ctrl-Z or enter end .	Use this mode to configure parameters for the switch interfaces. To configure multiple interfaces with the same parameters, see the "Configuring a Range of Interfaces" section on page 169.
Line configuration	While in global configuration mode, specify a line with the line vty command.	Switch(config-line)#	To exit to global configuration mode, enter exit . To return to privileged EXEC mode, press Ctrl-Z or enter end .	Use this mode to configure parameters for the terminal line.

Getting Help

You can enter a question mark (?) at the system prompt to display a list of commands available for each command mode. You can also obtain a list of associated keywords and arguments for any command, as shown in Table 2

Table 2. Help Summary

Command	Purpose
help	Obtain a brief description of the help system in any command mode.
<i>abbreviated-command-entry?</i>	Obtain a list of commands that begin with a particular character string. For example: Switch# di? dir disable disconnect
<i>abbreviated-command-entry<Tab></i>	Complete a partial command name. For example: Switch# sh conf<tab> Switch# show configuration
?	List all commands available for a particular command mode. For example: Switch> ?
<i>command ?</i>	List the associated keywords for a command. For example: Switch> show ?
<i>command keyword ?</i>	List the associated arguments for a keyword. For example: Switch(config)# cdp holdtime ? <10-255> Length of time (in sec) that receiver must keep this packet

Specifying Ports in Interface Configuration Mode

To configure a port, you need to specify the interface type, slot, and port number by using the **interface** configuration command. For example, to configure port 17 on a switch, you enter:

```
switch(config)# interface gi 0/17
```

- Interface type—Each switch platform supports different types of interfaces. To display a complete list of the interface types supported on your switch, enter the **interface ?** global configuration command.
- Slot number—The slot number on the switch. On the switch, the slot number is 0.
- Port number—The number of the physical port on the switch.

Switch ports 1 to 14 are internal 1000 Mbps connections to the other blades in the BladeCenter. These ports are configured to operate at 1000 Mbps in full-duplex mode.

Switch ports 15 and 16 are internal 100 Mbps connections to the Management Module. These ports are configured to operate at 100 Mbps in full-duplex mode.

Note: You cannot change the speed and duplex settings on the internal ports 1 to 16.

Switch ports 17 to 20 are for 10/100/1000 Mbps connections to external devices such as other switches. By default, these ports are configured for autonegotiation for speed and duplex mode. You can change these settings. You can change the speed and duplex settings on the external ports 17 to 20.

Note: The interface notation for switch ports 1 to 20 is **interface gigabitethernet** (such as **interface gi**).

Abbreviating Commands

You have to enter only enough characters for the switch to recognize the command as unique. This example shows how to enter the **show configuration** privileged EXEC command:

```
Switch# show conf
```

Using no and default Forms of Commands

Almost every configuration command also has a **no** form. In general, use the **no** form to disable a feature or function or reverse the action of a command. For example, the **no shutdown** interface configuration command reverses the shutdown of an interface. Use the command without the keyword **no** to re-enable a disabled feature or to enable a feature that is disabled by default.

Configuration commands can also have a **default** form. The **default** form of a command returns the command setting to its default. Most commands are disabled by default, so the **default** form is the same as the **no** form. However, some commands are enabled by default and have variables set to certain default values. In these cases, the **default** command enables the command and sets variables to their default values.

Understanding CLI Messages

Table 3 lists some error messages that you might encounter while using the CLI to configure your switch.

Table 3. Common CLI Error Messages

Error Message	Meaning	How to Get Help
% Ambiguous command: "show con"	You did not enter enough characters for your switch to recognize the command.	Re-enter the command followed by a question mark (?) with a space between the command and the question mark. The possible keywords that you can enter with the command are displayed.

Table 3. Common CLI Error Messages (continued)

Error Message	Meaning	How to Get Help
% Incomplete command.	You did not enter all the keywords or values required by this command.	Re-enter the command followed by a question mark (?) with a space between the command and the question mark. The possible keywords that you can enter with the command are displayed.
% Invalid input detected at '^' marker.	You entered the command incorrectly. The caret (^) marks the point of the error.	Enter a question mark (?) to display all the commands that are available in this command mode. The possible keywords that you can enter with the command are displayed.

Using Command History

The Cisco IOS provides a history or record of commands that you have entered. This feature is particularly useful for recalling long or complex commands or entries, including access lists. You can customize the command history feature to suit your needs as described in these sections:

- Changing the Command History Buffer Size, on page 29
- Recalling Commands, on page 29
- Disabling the Command History Feature, on page 30

Changing the Command History Buffer Size

By default, the switch records ten command lines in its history buffer. Beginning in privileged EXEC mode, enter this command to change the number of command lines that the switch records during the current terminal session:

```
Switch# terminal history [size number-of-lines]
```

The range is from 0 to 256.

Beginning in line configuration mode, enter this command to configure the number of command lines the switch records for all sessions on a particular line:

```
Switch(config-line)# history [size number-of-lines]
```

The range is from 0 to 256.

Recalling Commands

To recall commands from the history buffer, perform one of the actions listed in Table 4:

Table 4. Recalling Commands

Action ¹	Result
Press Ctrl-P or the up arrow key.	Recall commands in the history buffer, beginning with the most recent command. Repeat the key sequence to recall successively older commands.
Press Ctrl-N or the down arrow key.	Return to more recent commands in the history buffer after recalling commands with Ctrl-P or the up arrow key. Repeat the key sequence to recall successively more recent commands.
show history	While in privileged EXEC mode, list the last several commands that you just entered. The number of commands that are displayed is determined by the setting of the terminal history global configuration command and history line configuration command.

1.The arrow keys function only on ANSI-compatible terminals such as VT100s.

Disabling the Command History Feature

The command history feature is automatically enabled.

To disable the feature during the current terminal session, enter the **terminal no history** privileged EXEC command.

To disable command history for the line, enter the **no history** line configuration command.

Using Editing Features

This section describes the editing features that can help you manipulate the command line. It contains these sections:

- Enabling and Disabling Editing Features, on page 30
- Editing Commands through Keystrokes, on page 31
- Editing Command Lines that Wrap, on page 32

Enabling and Disabling Editing Features

Although enhanced editing mode is automatically enabled, you can disable it.

To re-enable the enhanced editing mode for the current terminal session, enter this command in privileged EXEC mode:

```
Switch# terminal editing
```

To reconfigure a specific line to have enhanced editing mode, enter this command in line configuration mode:

```
Switch(config-line)# editing
```

To globally disable enhanced editing mode, enter this command in line configuration mode:

```
Switch(config-line)# no editing
```


Editing Commands through Keystrokes

Table 5 shows the keystrokes that you need to edit command lines.

Table 5. *Editing Commands through Keystrokes*

Capability	Keystroke ¹	Purpose
Move around the command line to make changes or corrections.	Press Ctrl-B , or press the left arrow key.	Move the cursor back one character.
	Press Ctrl-F , or press the right arrow key.	Move the cursor forward one character.
	Press Ctrl-A .	Move the cursor to the beginning of the command line.
	Press Ctrl-E .	Move the cursor to the end of the command line.
	Press Esc B .	Move the cursor back one word.
	Press Esc F .	Move the cursor forward one word.
	Press Ctrl-T .	Transpose the character to the left of the cursor with the character located at the cursor.
Recall commands from the buffer and paste them in the command line. The switch provides a buffer with the last ten items that you deleted.	Press Ctrl-Y .	Recall the most recent entry in the buffer.
	Press Esc Y .	Recall the next buffer entry. The buffer contains only the last 10 items that you have deleted or cut. If you press Esc Y more than ten times, you cycle to the first buffer entry.
Delete entries if you make a mistake or change your mind.	Press the Delete or Backspace key.	Erase the character to the left of the cursor.
	Press Ctrl-D .	Delete the character at the cursor.
	Press Ctrl-K .	Delete all characters from the cursor to the end of the command line.
	Press Ctrl-U or Ctrl-X .	Delete all characters from the cursor to the beginning of the command line.
	Press Ctrl-W .	Delete the word to the left of the cursor.
	Press Esc D .	Delete from the cursor to the end of the word.
Capitalize or lowercase words or capitalize a set of letters.	Press Esc C .	Capitalize at the cursor.
	Press Esc L .	Change the word at the cursor to lowercase.
	Press Esc U .	Capitalize letters from the cursor to the end of the word.
Designate a particular keystroke as an executable command, perhaps as a shortcut.	Press Ctrl-V or Esc Q .	

Table 5. Editing Commands through Keystrokes (continued)

Capability	Keystroke ¹	Purpose
Scroll down a line or screen on displays that are longer than the terminal screen can display. Note: The More prompt is used for any output that has more lines than can be displayed on the terminal screen, including show command output. You can use the Return and Space bar keystrokes whenever you see the More prompt.	Press the Return key.	Scroll down one line.
	Press the Space bar.	Scroll down one screen.
Redisplay the current command line if the switch suddenly sends a message to your screen.	Press Ctrl-L or Ctrl-R .	Redisplay the current command line.

1.The arrow keys function only on ANSI-compatible terminals such as VT100s.

Editing Command Lines that Wrap

You can use a wraparound feature for commands that extend beyond a single line on the screen. When the cursor reaches the right margin, the command line shifts ten spaces to the left. You cannot see the first ten characters of the line, but you can scroll back and check the syntax at the beginning of the command.

To scroll back to the beginning of the command entry, press **Ctrl-B** or the left arrow key repeatedly. You can also press **Ctrl-A** to immediately move to the beginning of the line.

Note: The arrow keys function only on ANSI-compatible terminals such as VT100s.

In this example, the **access-list** global configuration command entry extends beyond one line. When the cursor first reaches the end of the line, the line is shifted ten spaces to the left and redisplayed. The dollar sign (\$) shows that the line has been scrolled to the left. Each time the cursor reaches the end of the line, the line is again shifted ten spaces to the left.

```
Switch(config)# access-list 101 permit tcp 131.108.2.5 255.255.255.0 131.108.1
Switch(config)# $ 101 permit tcp 131.108.2.5 255.255.255.0 131.108.1.20 255.25
Switch(config)# $t tcp 131.108.2.5 255.255.255.0 131.108.1.20 255.255.255.0 eq
Switch(config)# $108.2.5 255.255.255.0 131.108.1.20 255.255.255.0 eq 45
```

After you complete the entry, press **Ctrl-A** to check the complete syntax before pressing the **Return** key to execute the command. The dollar sign (\$) appears at the end of the line to show that the line has been scrolled to the right:

```
Switch(config)# access-list 101 permit tcp 131.108.2.5 255.255.255.0
131.108.1$
```

The software assumes you have a terminal screen that is 80 columns wide. If you have a width other than that, use the **terminal width** privileged EXEC command to set the width of your terminal.

Use line wrapping with the command history feature to recall and modify previous complex command entries. For information about recalling previous command entries, see the “Editing Commands through Keystrokes” section on page 31.

Searching and Filtering Output of `show` and `more` Commands

You can search and filter the output for **show** and **more** commands. This is useful when you need to sort through large amounts of output or if you want to exclude output that you do not need to see.

To use this functionality, enter a **show** or **more** command followed by the *pipe* character (`|`), one of the keywords **begin**, **include**, or **exclude**, and an expression that you want to search for or filter out:

```
command | {begin | include | exclude} regular-expression
```

Expressions are case sensitive. For example, if you enter `| exclude output`, the lines that contain *output* are not displayed, but the lines that contain *Output* are displayed.

This example shows how to include in the output display only lines where the expression *protocol* appears:

```
Switch# show interfaces | include protocol
Vlan1 is up, line protocol is up
Vlan10 is up, line protocol is down
GigabitEthernet0/17 is up, line protocol is down
GigabitEthernet0/20 is up, line protocol is up
```

Accessing the CLI

You can access the CLI through a remote Telnet session, but your switch must first be configured for this type of access. For more information, see the “Changing a Telnet Password for a Terminal Line” section on page 121.

You can establish a connection with the switch by using any Telnet TCP/IP or encrypted Secure Shell (SSH) package from a remote management station. The switch must have network connectivity with the Telnet or SSH client, and the switch must have an enable secret password configured.

For information about configuring the switch for Telnet access, see the “Changing a Telnet Password for a Terminal Line” section on page 121. The switch supports up to 16 simultaneous Telnet sessions. Changes made by one Telnet user are reflected in all other Telnet sessions.

For information about configuring the switch for SSH, see the “Configuring the Switch for Secure Shell” section on page 146. The switch supports up to five simultaneous secure SSH sessions.

After you connect through a Telnet session or through an SSH session, the user EXEC prompt appears on the management station.

Accessing the CLI from a Browser

This procedure assumes that you have met the software requirements (including browser and Java plug-in configurations) and have assigned IP information and a Telnet password to the switch or command switch, as described in the release notes.

To access the CLI from a web browser, follow these steps:

1. Start one of the supported browsers.
2. In the **URL** field, enter the IP address of the command switch.
3. When the Cisco Systems Access page appears, click **Telnet** to start a Telnet session.
4. Enter the switch password.

The user EXEC prompt appears on the management station.

Note: Copies of the CMS pages that you display are saved in your browser memory cache until you exit the browser session. A password is not required to redisplay these pages, including the Cisco Systems Access page. You can access the CLI by clicking **Web Console - HTML access to the command line interface** from a cached copy of the Cisco Systems Access page. To prevent unauthorized access to CMS and the CLI, exit your browser to end the browser session.

Chapter 3. Getting Started with CMS

This chapter describes the Cluster Management Suite (CMS) on the switch. Be sure to also see the release notes before proceeding with this chapter. It contains these topics:

- Launching CMS, on page 35
- Features, on page 36
- Front Panel View, on page 38
- Topology View, on page 42
- Menus and Toolbar, on page 45
- Interaction Modes, on page 54
- CMS Window Components, on page 56
- Accessing CMS, on page 58
- Saving Your Configuration, on page 60
- Restoring Your Configuration, on page 61
- CMS Preferences, on page 61
- Using Different Versions of CMS, on page 61
- Where to Go Next, on page 61

It does not contain:

- Procedures for using the configuration windows in CMS. The online help gives this information.
- System requirements and procedures for browser and Java plug-in configuration. The release notes give this information.

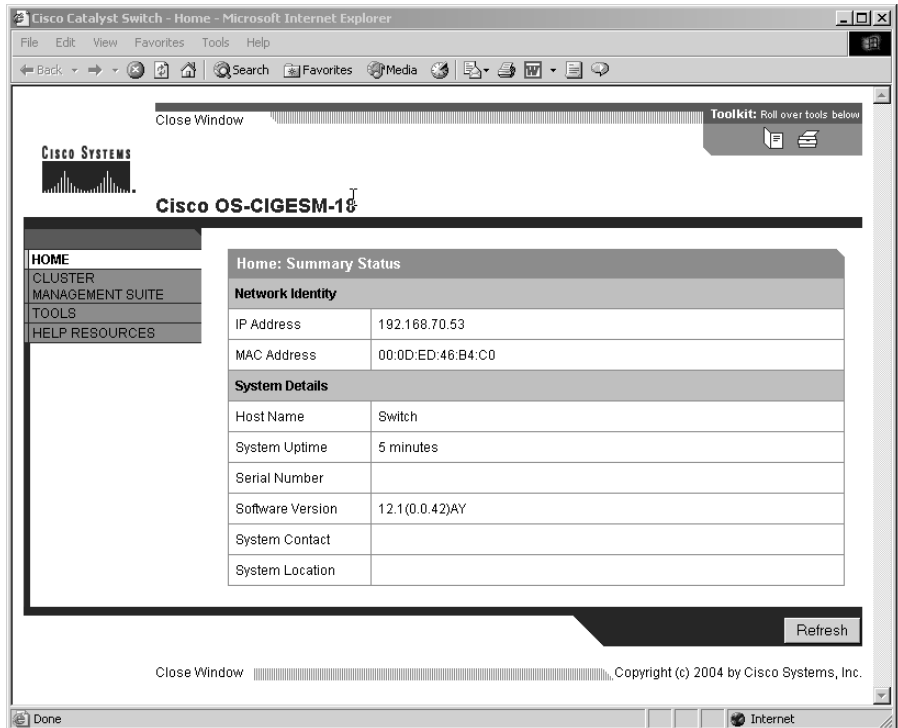
Refer to the appropriate switch documentation for descriptions of the web-based management software used on other switches.

Launching CMS

Before you launch CMS, follow the steps for setting up your switch and assigning it an IP address.

After you assign the IP address, enter the IP address of your switch and your password (if one has been set) in a browser window. Then the Switch Home Page opens (see Figure 4). To launch CMS, click **Cluster Management Suite** on the left side of the Switch Home Page.

Figure 4. Switch Home Page



The Switch Home Page has these tabs:

- Cluster Management Suite—Launches CMS, through which you can manage the switch
- Tools—Accesses diagnostic and monitoring tools
- Help Resources—Provides links to the IBM website and technical documentation

Features

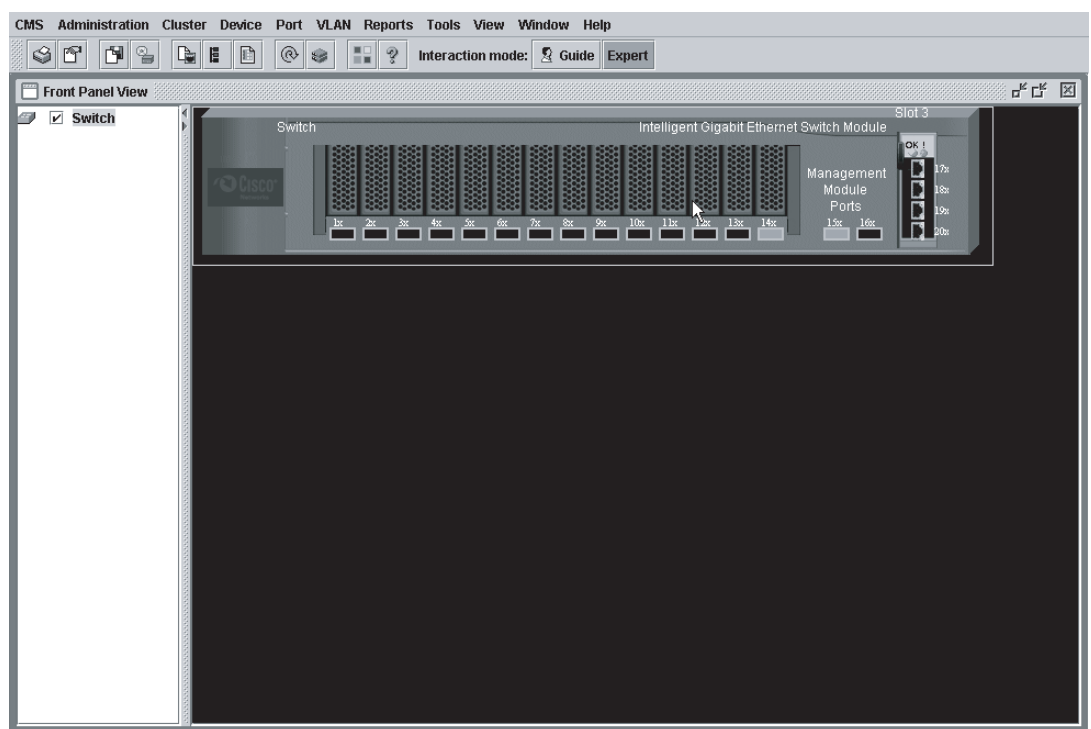
CMS provides these features for managing switch clusters and individual switches from Web browsers such as Netscape Communicator or Microsoft Internet Explorer:

- Two views of your network, as shown in Figure 5, that can be displayed at the same time:
 - A Front Panel view that displays the front-panel image of a specific set of switches in a cluster. From this view, you can select multiple ports or multiple switches and configure them with the same settings.

Note: When CMS is launched from a command switch, the Front Panel view displays the front-panel image of the command switch. You can select more switches to be displayed. When CMS is launched from a noncommand switch, the Front Panel view displays only the front panel of the specific switch.
 - A Topology view that displays a network map that uses icons representing switch clusters, the command switch, cluster members, cluster candidates, neighboring devices that are not eligible to join a cluster, and link types. From this view, you can select multiple switches and configure them to run with the same settings. You can also display link information in the form of link reports and link graphs.

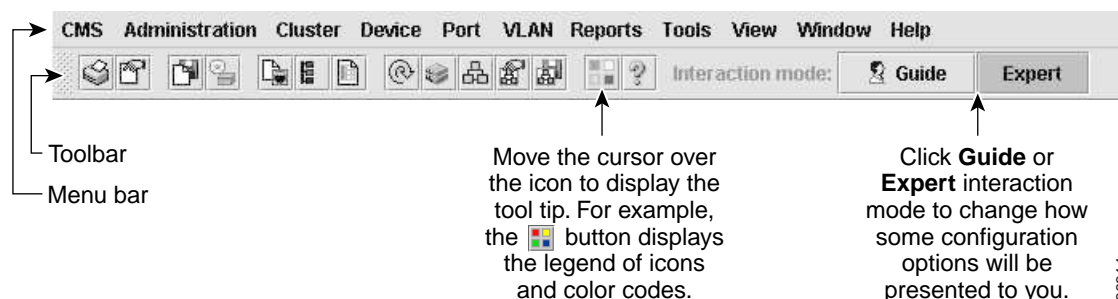
This view is available only when CMS is launched from a command switch.

Figure 5. CMS Front Panel



- Menus and a toolbar, as shown in Figure 6, to access configuration and management options:
 - The menu bar provides the complete list of options for managing a single switch and switch clusters.
 - The toolbar provides buttons for commonly used switch and cluster configuration options and information windows such as legends and online help.
 - The port popup menu, in the Front Panel view, provides options specific for configuring and monitoring switch ports.
 - The device popup menu, in either the Front Panel or the Topology view, provides switch and cluster configuration and monitoring options.
 - The candidate, member, and link popup menus provide options for configuring and monitoring devices and links in the Topology view.

Figure 6. CMS Menus and Toolbar




- Tools to simplify configuration tasks:
 - Interactive modes—guide mode and expert mode—that control the presentation of some complex configuration options.

- Wizards that require minimal information from you to configure some complex features.
- Comprehensive online help that gives high-level concepts and procedures for performing CMS tasks.
- Two levels of access to the configuration options: read-write access for users allowed to change switch settings and read-only access for users allowed to only view switch settings.
- Consistent set of GUI components (such as tabs, buttons, drop-down lists, and tables) for a uniform approach to viewing and setting configuration parameters.

Front Panel View

When CMS is launched from a command switch, the Front Panel view displays the front-panel image of the command switch, as shown in Figure 7. You can select switches to be displayed by checking the boxes in the cluster tree view (left panel of CMS). The switches that are displayed in the tree view can be re-arranged by dragging and dropping them.

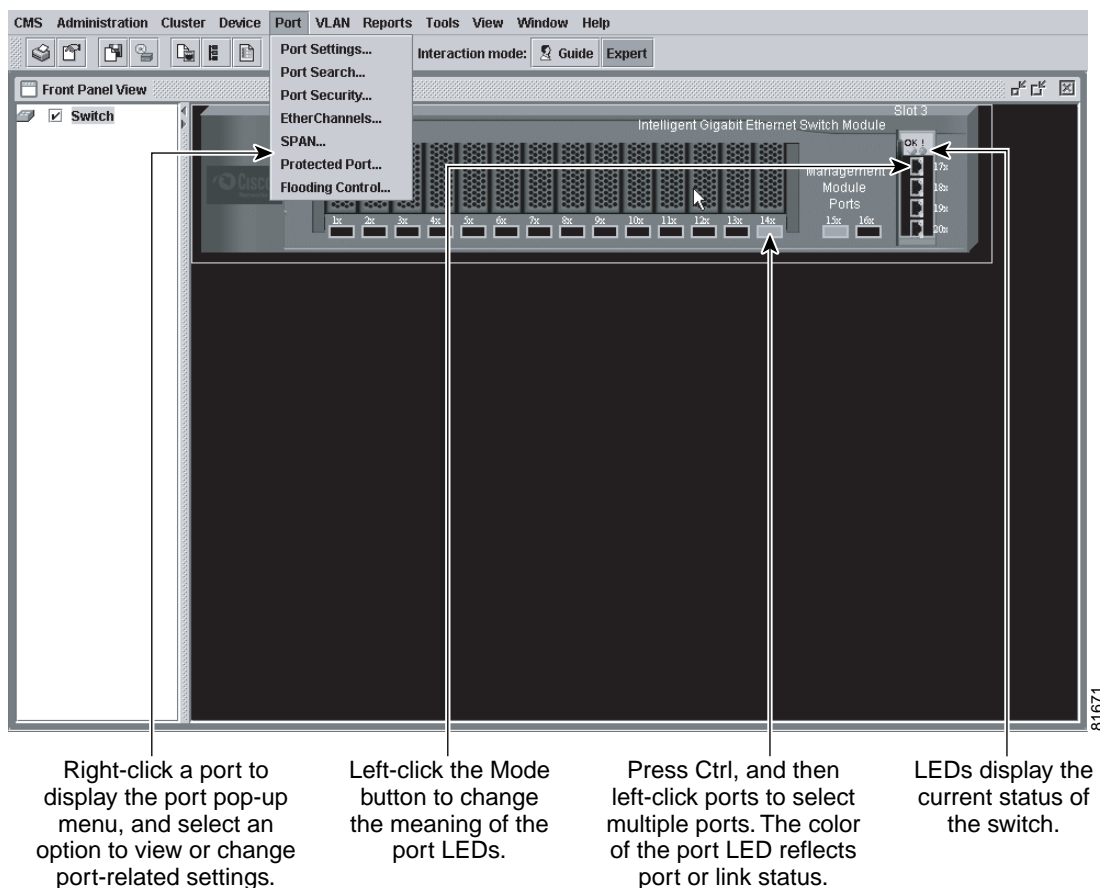
Figure 7. Front Panel View from a Command Switch



Note: CMS from a standalone switch or from a noncommand switch is referred to as Device Manager (also referred to as Switch Manager). Device Manager is for configuring an individual switch. When you select Device Manager for a specific switch in the cluster, you launch a separate CMS session. The Device Manager interface can vary between the switch platforms.

When CMS is launched from a standalone or noncommand member switch, the Front Panel view displays only the front panel of the specific switch, as shown in Figure 8.

Figure 8. Front Panel View from a Standalone Switch



Right-click a port to display the port pop-up menu, and select an option to view or change port-related settings.

Left-click the Mode button to change the meaning of the port LEDs.

Press Ctrl, and then left-click ports to select multiple ports. The color of the port LED reflects port or link status.

LEDs display the current status of the switch.

Cluster Tree

Figure 7 shows the cluster tree that appears in the left frame of the Front Panel view and shows the name of the cluster and a list of its members. Figure 9 shows the device icons that you can drag and drop to rearrange them in the cluster tree. The colors of the devices in the cluster tree show the status of the devices, as listed in Table 6.

If you want to configure switch or cluster settings on one or more switches, select the appropriate front-panel images.

- To select a front-panel image, click either the cluster-tree icon or the corresponding front-panel image. The front-panel image is then highlighted with a yellow outline.
- To select multiple front-panel images, press the **Ctrl** key, and left-click the cluster-tree icons or the front-panel images. To deselect an icon or image, press the **Ctrl** key, and left-click the icon or image.

If the cluster has many switches, you might need to scroll down the window to display the rest of the front-panel images. Instead of scrolling, you can click an icon in the cluster tree, and CMS then scrolls and displays the corresponding front-panel image.

Figure 9. Cluster-Tree Icons



Table 6. Cluster Tree Icon Colors

Color	Device Status
Green	Switch is operating normally.
Red	Switch is not powered on or has lost power, or the command switch is unable to communicate with the member switch.

Front-Panel Images

You can manage the switch from a remote station by using the front-panel images. The front-panel images are updated based on the network polling interval that you set from **CMS > Preferences**.


This section includes descriptions of the LED images. Similar descriptions of the switch LEDs are provided in the switch hardware installation guide.

Note: The Preferences window is available if your switch access level is read-only. For more information about the read-only access mode, see the “Access Modes in CMS” section on page 59.

Figure 10 shows the port icons as they appear in the front-panel images. To select a port, click the port on the Front Panel view. The port is then highlighted with a yellow outline. To select multiple ports, you can:

- Press the left mouse button, drag the pointer over the group of ports that you want to select, and then release the mouse button.
- Press the Ctrl key, and click the ports that you want to select.
- Right-click a port, and select **Select All Ports** from the port popup menu.

Figure 10. Port Icons



Port Modes and LEDs

Table 7 lists the port modes that determine the type of information displayed through the port LEDs. When you change port modes, the meanings of the port LED colors also change.

Note: The bandwidth utilization mode (UTIL LED) does not appear on the front-panel images. Select **Reports > Bandwidth Graphs** to display the total bandwidth in use by the switch. Refer to the switch hardware installation guide for information about using the UTIL LED.

Table 7. Port LEDs

Indicator Name	Port LED Color	Description
OK	Green	Power is on.
Fault	Amber	Switch failure to determine the fault. Go to the Management module system status web screen and select the appropriate bay.
Link OK	Green (off)	No link present.
	Green	Link present.
Link OK	Green (off)	No traffic
	Green (flashing)	Link activity

VLAN Membership Modes

Table 8 lists the colors that outline the ports (Front Panel view) when you click **Highlight VLAN Port Membership Modes** on the Configure VLANs tab on the VLAN window. The colors show the VLAN membership mode of each port. The VLAN

membership mode determines the kind of traffic the port carries and the number of VLANs to which it can belong. For more information about these modes, see the “VLAN Port Membership Modes” section on page 244.

Table 8. VLAN Membership Modes

Mode	Color
Static access	Light green
Dynamic access	Pink
802.1Q trunk	Peach
Negotiate trunk	White

Topology View

The Topology view displays how the devices within a switch cluster are connected and how the switch cluster is connected to other clusters and devices. From this view, you can add and remove cluster members. This view provides two levels of detail of the network topology:

- **Expand Cluster:** When you right-click a cluster icon and select Expand Cluster, the Topology view displays the switch cluster in detail, as shown in Figure 11. This view shows the command switch and member switches in a cluster. It also shows candidate switches that can join the cluster. This view does not display the details of any neighboring switch clusters
- **Collapse Cluster:** When you right-click a command-switch icon and select Collapse Cluster, the cluster is collapsed and represented by a single icon, as shown in Figure 12. The view shows how the cluster is connected to other clusters, candidate switches, and devices that are not eligible to join the cluster (such as routers, access points, IP phones, and so on).

Note: The Topology view displays only the switch cluster and network neighborhood of the specific command or member switch that you access. To display a different switch cluster, you need to access the command switch or member switch of that cluster.

You can arrange the device icons in either view. To move a device icon, click and drag the icon. To select multiple device icons, you can either:

- Press the left mouse button, drag the pointer over the group of device icons that you want to select, and then release the mouse button.
- Press the Ctrl key, and click the device icons that you want to select.

After selecting the icons, drag the icons to any area in the view.

Figure 11. Expand Cluster View

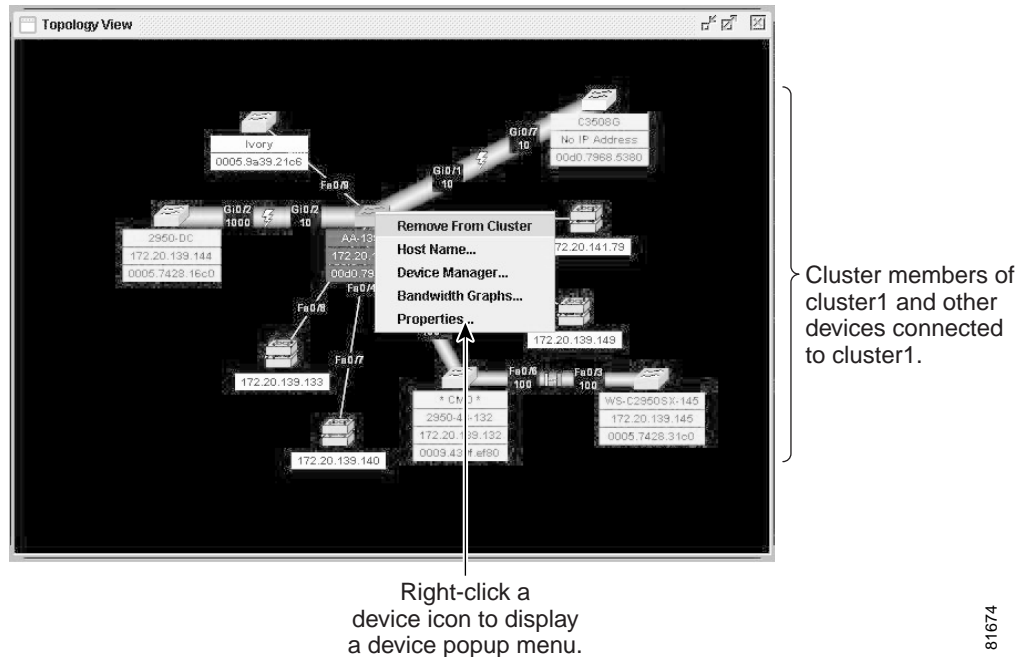
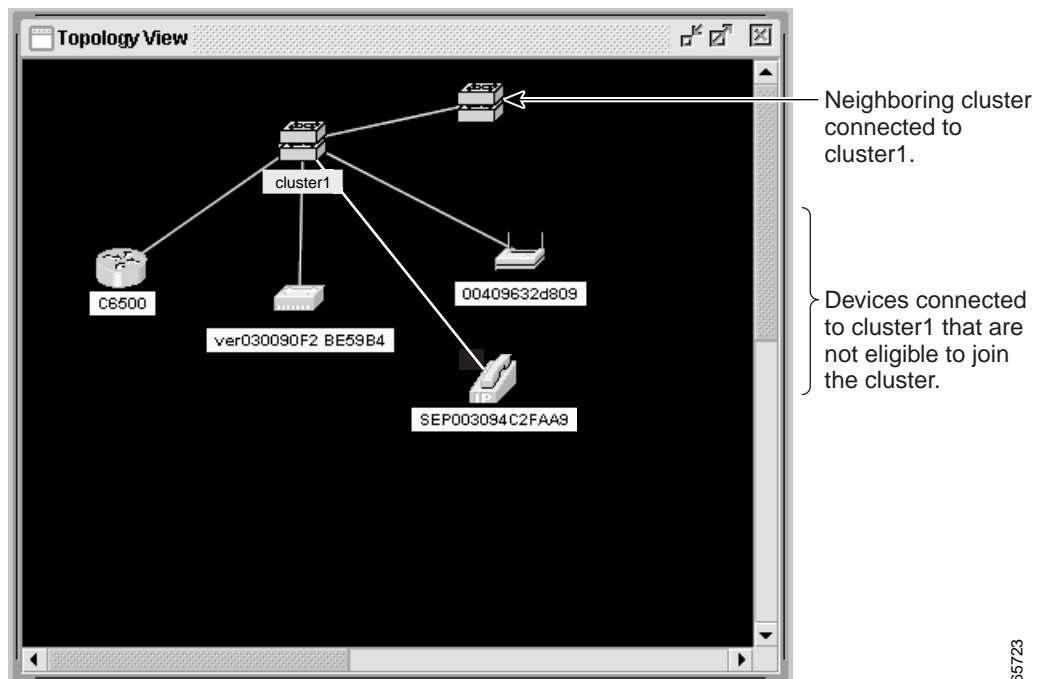


Figure 12. Collapse Cluster View



Topology Icons and Labels

The Topology view and the cluster tree use the same set of device icons to represent clusters, command and standby command switches, and member switches. They also use the same labels to identify the command switch (**CMD**) and the standby command switch (**STBY**).

The Topology view uses additional icons to represent these types of neighboring devices:

- Devices that are not eligible to join the cluster, such as Cisco IP Phones, Cisco access points, and Cisco Discovery Protocol (CDP)-capable hubs and routers
- Devices that are identified as unknown devices, such as some Cisco devices and third-party devices

Important: Neighboring devices are only displayed if they are connected to cluster members. To display neighboring devices in the Topology view, either add the switch to which they are connected to a cluster, or enable that switch as a command switch.

Note: Candidate switches are distinguished by the color of their device label. Device labels and their colors are described in the “Colors in the Topology View” section on page 44.

To select a device, click the icon. The icon is then highlighted. To select multiple devices, you can either:

- Press the left mouse button, drag the pointer over the group of icons that you want to select, and then release the mouse button.
- Press the **Ctrl** key, and click the icons that you want to select.

The Topology view also uses a set of link icons to show the link type and status between two devices. To select a link, click the link that you want to select. To select multiple links, press the Ctrl key, and click the links that you want to select.

Device and Link Labels

The Topology view displays device and link information by using these labels:

- Cluster and switch names
- Switch MAC and IP addresses
- Link type between the devices
- Link speed and IDs of the interfaces on both ends of the link

When using these labels, keep these considerations in mind:

- The IP address displays only in the labels for the command switch and member switches.
- The label of a neighboring cluster icon only displays the IP address of the command-switch IP address.
- The link speeds displayed are the actual link speeds.

You can change the label settings from the Topology Options window by selecting **View > Topology Options**.

Colors in the Topology View

The colors of the Topology view icons show the status of the devices and links, as listed in Table 9, Table 10, and Table 11.

Table 9. Device Icon Colors

Icon Color	Color Meaning
Green	The device is operating.
Red ¹	The device is not operating.

Table 10. Single Link Icon Colors

Link Color	Color Meaning
Green	Active link
Red	Down or blocked link

Table 11. Multiple Link Icon Colors

Link Color	Color Meaning
Both green	All links are active.
One green; one red	At least one link is active, and at least one other link is down or blocked.
Both red	All links are down or blocked.

The color of a device label shows the cluster membership of the device, as listed in Table 12.

Table 12. Device Label Colors

Label Color	Color Meaning
Green	A cluster member, either a member switch or the command switch
Cyan	A candidate switch that is eligible to join the cluster
Yellow	An unknown device or a device that is not eligible to join the cluster

Topology Display Options

You can set the type of information displayed in the Topology view by changing the settings in the Topology Options window. To display this window, select **View > Topology Options**. From this window, you can select:

- Device icons (including IP Phones, CPEs, Neighbors, Access Points, and Candidates) that you want displayed in or filtered from the Topology View window
- Interface IDs and Actual Speed values that you want displayed in the Link window
- Host Names, IP addresses, and MAC address labels that you want displayed in the Node window

Menus and Toolbar

The configuration and monitoring options for configuring switches and switch clusters are available from menus and a toolbar.

Menu Bar

The menu bar, as shown in Figure 6, provides the complete list of options for managing a single switch and switch cluster.

Options displayed from the menu bar can vary:

- The option for enabling a command switch is only available from a CMS session launched from a command-capable switch.
- Cluster management tasks, such as upgrading the software of groups of switches, are available only from a CMS session launched from a command switch.

- If you launch CMS from a specific switch, the menu bar displays the features supported only by that switch.
- If you launch CMS from a command switch, the menu bar displays the features supported on the switches in the cluster, with one exception. If the command switch is a Layer 2 switch, such as a Catalyst 2950 or Catalyst 3500 XL switch, the menu bar displays the features of all Layer 2 switches in the cluster.
- We strongly recommend that the highest-end, command-capable switch in the cluster be the command switch. If your switch cluster has Catalyst 2900 XL, Catalyst 2940, Catalyst 2950, Catalyst 2955, and Catalyst 3500 XL switches, the Catalyst 2950 or Catalyst 2955 switch should be the command switch.

Note: The IGESM should never be the command switch.

- Standby command switches must meet these requirements:
 - When the command switch is a Catalyst 2955 switch, all standby command switches must be Catalyst 2955 switches.
 - When the command switch is a Catalyst 2950 LRE switch, all standby command switches must be Catalyst 2950 LRE switches.
 - When the command switch is a non-LRE Catalyst 2950 switch running Cisco IOS Release 12.1(9)EA1 or later, all standby command switches must be non-LRE Catalyst 2950 switches running Cisco IOS Release 12.1(9)EA1 or later.
 - When the command switch is a non-LRE Catalyst 2950 switch running Cisco IOS Release 12.1(6)EA2 or later, all standby command switches must be non-LRE Catalyst 2950 switches running Cisco IOS Release 12.1(6)EA2 or later.
 - When the command switch is running Cisco IOS Release 12.0(5)WC2 or earlier, the standby command switches can be these switches: Catalyst 2900 XL, non-LRE Catalyst 2950, and Catalyst 3500 XL switches.

It is preferable that the command switch and standby command switches are of the same switch platform and that both are running the same level of software. In the event of a failover, the standby command switch must support the same configuration and services that are running on the command switch.

Refer to the release notes for the Catalyst switches that can be part of a switch cluster.

Unless noted otherwise, the menu-bar options in the list that follows are available from a Catalyst 2950 or Catalyst 2955 command switch when the cluster contains only Catalyst 2950 or Catalyst 2955 member switches. The menu bar of the command switch displays all menu-bar options available from the cluster, including options from member switches from other cluster-capable switch platforms.

Note: Access modes affect the availability of features from CMS. Some CMS features are not available in read-only mode. For more information about how access modes affect CMS, see the “Access Modes in CMS” section on page 59.

These are the menu bar options:

- **CMS**
 - **Page Setup** - Set default document printer properties to be used when printing from CMS.
 - **Print Preview** - View the way the CMS window or help file will appear when printed.
 - **Print** - Print a CMS window or help file.

- **Guide Mode/Expert Mode** - Select which interaction mode to use when you select a configuration option (not available in read-only mode).
- **Preferences** - Set CMS display properties, such as polling intervals, the default views to open at startup, and the color of administratively shutdown ports. Some options from this menu are not available in read-only mode.
- **Administration**
 - **IP Addresses** - Configure IP information for a switch. Some options from this menu are not available in read-only mode.
 - **SNMP** - Enable and disable Simple Network Management Protocol (SNMP), enter community strings, and configure end stations as trap managers. Some options from this menu are not available in read-only mode.
 - **System Time** - Configure the system time or configure the Network Time Protocol (NTP). Some options from this menu are not available in read-only mode.
 - **HTTP Port** - Configure the Hypertext Transfer Protocol (HTTP) port number. Some options from this menu are not available in read-only mode.
 - **Users and Passwords** - Configure usernames and passwords for privilege levels 0 to 15.
 - **Console Baud Rate** - Change the baud rate for the switch service port. Some options from this menu are not available in read-only mode.
 - **MAC Addresses** - Enter dynamic, secure, and static addresses in a switch address table. You can also define the forwarding behavior of static addresses. Some options from this menu are not available in read-only mode.
 - **ARP** - Display the device Address Resolution Protocol (ARP) table, and configure the ARP cache timeout setting. Some options from this menu are not available in read-only mode.
 - **Save Configuration** - Save the configuration for the cluster or switch to flash memory (not available in read-only mode).
 - **Restore Configuration** - Restore the configuration file to one or more switches in the cluster.
 - **Software Upgrade** - Upgrade the software for the cluster or a switch (not available in read-only mode).
 - **System Reload** - Reboot the switch with the latest installed software (not available in read-only mode).
 - **Event Notification** - Create notification IDs that generate e-mail notifications when system events occur.
- **Cluster**
 - **Cluster Manager** - Launch a CMS session from the member switch (available only from a Device Manager session on a cluster member).
 - **Create Cluster** - Designate a command switch, and name a cluster (not available in read-only mode). This option is available only from a Device Manager session on a command-capable switch that is not a cluster member.
 - **Delete Cluster** - Delete a cluster (not available in read-only mode). This option is available only from a cluster management session.
 - **Add to Cluster** - Add a candidate to a cluster (not available in read-only mode). This option is available only from a cluster management session.
 - **Remove from Cluster** - Remove a member from the cluster (not available in read-only mode) This option is available only from a cluster management session.

- **Standby Command Switches** - Create a Hot Standby Router Protocol (HSRP) standby group to provide command-switch redundancy. Some options from this menu are not available in read-only mode. This option is available only from a cluster management session.
- **Hop Count** - Enter the number of hops away that a command switch looks for members and for candidate switches. Some options from this menu are not available in read-only mode. This option is available only from a cluster management session.
- **Device**
 - **Device Manager** - Launch Device Manager for a specific switch.
 - **Host Name** - Change the host name of a switch (not available in read-only mode).
 - **STP** - Display and configure STP parameters for a switch. Some options from this menu are not available in read-only mode.
 - **IGMP Snooping** - Enable and disable Internet Group Management Protocol (IGMP) snooping and IGMP Immediate-Leave processing on the switch. Join or leave multicast groups, and configure multicast routers. Some options from this menu are not available in read-only mode.
 - **ACL** (guide mode available in read-write mode) - Create and maintain access control lists (ACLs), and attach ACLs to specific ports. Some options from this menu are not available in read-only mode.
 - **Security Wizard** - Filter certain traffic, such as HTTP traffic, to certain networks or devices. Restrict access to servers, networks, or application data from certain networks or devices (not available in read-only mode).
 - **QoS** - Display submenu options to configure, enable, and disable quality of service (QoS) parameters for Trust settings, Queues, Maps, Classes (guide mode available), and Policies (guide mode available). Some options from this menu are not available in read-only mode.
 - **AVVID Wizards** - Configure a port to send or receive voice traffic by using the Voice Wizard. Optimize multiple video servers for sending video traffic by using the Video Wizard. Provide a higher priority to specific applications by using the Data Wizard.
- Note:** AVVID Wizards are not available in read-only mode.
- **Port**
 - **Port Settings** - Display and configure port parameters on a switch. Some options from this menu are not available in read-only mode.
 - **Port Search** - Search for a port through its description.
 - **Port Security** - Enable port security on a port (not available in read-only mode).
 - **EtherChannels** - Group ports into logical units for high-speed links between switches. Some options from this menu are not available in read-only mode.
 - **SPAN** - Enable Switched Port Analyzer (SPAN) port monitoring. Some options from this menu are not available in read-only mode.
 - **Protected Port** - Configure a port to prevent it from receiving bridged traffic from another port on the same switch. Some options from this menu are not available in read-only mode.
 - **Flooding Control** - Block the normal flooding of unicast and multicast packets and enable the switch to block packet storms. Some options from this menu are not available in read-only mode.
- **VLAN**















- **VLAN** (guide mode available in read-write mode) - Display VLAN membership, assign ports to VLANs, and configure 802.1Q trunks. Display and configure the VLAN Trunking Protocol (VTP) for interswitch VLAN membership. Some options from this menu are not available in read-only mode.
- **Management VLAN** - Change the management VLAN on the switch. Some options from this menu are not available in read-only mode.
- **VMPS** - Configure the VLAN Membership Policy Server (VMPS). Some options from this menu are not available in read-only mode.
- **Voice VLAN** - Configure a port to use a voice VLAN for voice traffic, separating it from the VLANs for data traffic. Some options from this menu are not available in read-only mode.
- **Reports**
 - **Inventory** - Display the device type, software version, IP address, and other information about a switch.
 - **Port Statistics** - Display port statistics.
 - **Bandwidth Graphs** - Display graphs that plot the total bandwidth in use by the switch.
 - **Link Graphs** - Display a graph showing the bandwidth being used for the selected link.
 - **Link Reports** - Display the link report for two connected devices. If one device is an unknown device or a candidate, only the cluster-member side of the link appears.
 - **ACL Reports** - Display a report about ACL statistics.
 - **Multicast** - Display a submenu to run an IGMP report.
 - **Resource Monitor** - Display masks for ACL and QoS policy maps.
 - **System Messages** - Display the most recent system messages (Cisco IOS messages and switch-specific messages) sent by the switch software. For more information about system messages, refer to the switch system message guide for that release.
- **Tools**
 - **Ping and Trace** - Perform a ping or Layer 2 traceroute operation on or to a specific address.
- **View**
 - **Refresh** - Update the views with the latest status.
 - **Front Panel** - Display the Front Panel view.
 - **Topology** - Display the Topology view. This option is available only from a cluster management session.
 - **Topology Options** - Select the information to be displayed in the Topology view.
 - **Automatic Topology Layout** - Request CMS to rearrange the topology layout. This option is available only from a cluster management session.
 - **Save Topology Layout** - Save the presentation of the cluster icons that you arranged in the Topology view to flash memory (not available in read-only mode). This option is available only from a cluster-management session.
- **Window**—List the open windows in your CMS session.
- **Help**
 - **Overview** - Obtain an overview of the CMS interface.

- **What's New** - Obtain a description of the new CMS features.
- **Help For Active Window** - Display the help for the active open window. You can also click **Help** from the active window.
- **Contents** - List all of the available online help topics.
- **Legend** - Display the legend that describes the icons, labels, and links.
- **About** - Display the CMS version number.

Toolbar

The toolbar buttons display commonly-used switch and cluster configuration options and information windows such as legends and online help. Hover the cursor over an icon to display the feature. Table 13 lists the toolbar options from left to right on the toolbar.

Table 13. Toolbar Buttons

Toolbar Option	Icon	Keyboard Shortcut	Task
Print		Ctrl-P	Print a CMS window or help file.
Preferences ¹		Ctrl-R	Set CMS display properties, such as polling intervals, the views to open at CMS startup, and the color of administratively shutdown ports.
Save Configuration ²		Ctrl-S	Save the configuration of the cluster of a switch to flash memory.
Software Upgrade ²		Ctrl-U	Upgrade the software for the cluster or a switch.
Port Settings ¹		—	Display and configure port parameters on a switch.
VLAN ¹		Ctrl-V	Display VLAN membership, assign ports to VLANs, and change the administration mode.
Inventory		Ctrl-T	Display the device type, the software version, the IP address, and other information about a switch.
Refresh		—	Update the views with the latest status.
Front Panel		—	Display the Front Panel view.
Topology ³		—	Display the Topology view.
Topology Options ³		—	Select the information to be displayed in the Topology view.
Save Topology ^{2 3} Layout		—	Save the presentation of the cluster icons that you arranged in the Topology view to flash memory.
Legend		—	Display the legend that describes the icons, labels, and links.
Help for Active Window		F1 key	Display the help for the active open window. You can also click Help from the active window.

1. Not available in read-only mode. For more information about the read-only and read-write access modes, see the “Access Modes in CMS” section on page 59.
2. Some options from this menu option are not available in read-only mode.
3. Available only from a cluster-management session.

Front Panel View Popup Menus

These popup menus are available in the Front Panel view:

Device Popup Menu

You can display all switch and cluster configuration windows from the menu bar, or you can display commonly-used configuration windows from the device popup menu, as listed in Table 14. To display the device popup menu, click the switch icon from the cluster tree or the front-panel image itself, and right-click.

Table 14. Device Popup Menu

Popup Menu Option	Task
Device Manager ¹	Launch Device Manager for the switch.
Host Name ²	Change the name of the switch.
Delete Cluster ^{2 3 4}	Delete a cluster.
Remove from Cluster ^{2 4}	Remove a member from the cluster.
Bandwidth Graphs	Display graphs that plot the total bandwidth in use.
Properties	Display information about the device and port on either end of the link and the state of the link.

1. Available from a cluster member switch but not from the command switch.
2. Not available in read-only mode. For more information about the read-only mode, see the “Access Modes in CMS” section on page 59.
3. Available only from the command switch.
4. Available only from a cluster-management session.

Port Popup Menu

You can display all port configuration windows from the Port menu on the menu bar, or you can display commonly-used port configuration windows from the port popup menu, as listed in Table 15. To display the port popup menu, click a specific port image, and right-click.

Table 15. Port Popup Menu

Popup Menu Option	Task
Port Settings ¹	Display and configure port settings.
VLAN ¹	Define the VLAN mode for a port or ports and add ports to VLANs.
Port Security ^{1 2 3}	Enable port security on a port.
Link Graphs ⁴	Display a graph showing the bandwidth used by the selected link.
Select All Ports	Select all ports on the switch for global configuration.

1. Some options from this menu are not available in read-only mode.
2. Available on switches that support the Port Security feature.
3. This feature is not available in read-only mode.
4. Available only when there is an active link on the port (that is, the port LED is green when in port status mode).

Topology View Popup Menus

These popup menus are available in the Topology view.

Link Popup Menu

Table 16 lists the reports and graphs that you can display for a specific link in the Topology view. To display the link popup menu, click the link icon, and right-click.

Table 16. Link Popup Menu


Link Popup Menu	Task
Link Report	Display the link report for two connected devices. If one device is an unknown device or a candidate, only the cluster member side of the link displays.
Link Graph	Display a graph showing the current bandwidth used by the selected link. You can change the graph polling interval by selecting CMS > Preferences .
Properties	Display information about the device and port on either end of the link and the state of the link.

The Link Report and Link Graph options are not available if at both ends of the link are one of these:

- Candidate switches
- Devices that are not eligible to join the cluster

If multiple links are configured between two devices, when you click the link icon and right-click, the Logical Link Content window appears, as shown in Figure 13. Click the link icon in this window, and right-click to display the link popup menu specific for that link.

Figure 13. Logical Link Window



Device Popup Menus

Table 17 through Table 22 list the popup menus for specific devices:

- Cluster (Table 17)
- Command switch (Table 18)
- Member or standby command switch (Table 19)
- Candidate switch with an IP address (Table 20)
- Candidate switch without an IP address (Table 21)
- Neighboring devices (Table 22)

To display a device popup menu, click an icon, and right-click.

Table 17. Cluster Icon Popup Menu

Popup Menu Option	Task
Expand cluster	View a cluster-specific topology view.
Properties	Display information about the device.

Table 18. Command-Switch Icon Popup Menu

Popup Menu Option	Task
Collapse cluster	View the neighborhood outside a specific cluster.
Host Name	Change the host name of a switch.
Bandwidth Graphs	Display graphs that plot the total bandwidth in use by the switch.
Properties	Display information about the device.

Table 19. Member or Standby Command-Switch Icon Popup Menu

Popup Menu Option	Task
Remove from Cluster	Remove a member from the cluster.
Host Name1	Change the host name of a switch.
Device Manager	Launch Device Manager for a switch.
Bandwidth Graphs	Display graphs that plot the total bandwidth in use by the switch.
Properties	Display information about the device.

Table 20. Candidate-Switch Icon Popup Menu (When the Candidate Switch Has an IP Address)

Popup Menu Option	Task
Add to Cluster	Add a candidate to a cluster.
Device Manager	Launch Device Manager for a switch.
Properties	Display information about the device.

Table 21. Candidate-Switch Icon Popup Menu (When the Candidate Switch Does Not Have an IP Address)

Popup Menu Option	Task
Add to Cluster	Add a candidate to a cluster.
Properties	Display information about the device.

Table 22. Neighboring-Device Icon Popup Menu

Popup Menu Option	Task
Device Manager	Access the web management interface of the device. Note: This option is available on Cisco access points, but not on Cisco IP Phones, hubs, routers and on unknown devices such as some Cisco devices and third-party devices.
Disqualification Code	Display the reason why the device could not join the cluster.
Properties	Display information about the device.

Interaction Modes

You can change the interaction mode of CMS to either guide or expert mode. Guide mode steps you through each feature option and provides information about the parameter. Expert mode displays a configuration window in which you configure the feature options.

Note: You cannot switch modes for an open CMS window (for example, from Guide Mode to Expert Mode). For the mode change to take effect on any other open CMS window, you need to close that window and then re-open it after you select the new mode.

Guide Mode

Guide mode is for users who want a step-by-step approach for completing a specific configuration task. This mode is not available for all features. A menu-bar option that has a person icon means that guide mode is available for that option.

When you click **Guide Mode** and then select a menu-bar option that supports it, CMS displays a specific parameter of the feature with information about the parameter field. To configure the feature, you provide the information that CMS requests in each step until you click **Finish** in the last step. Clicking **Cancel** at any time closes and ends the configuration task without applying any changes.

If **Expert Mode** is selected and you want to use Guide Mode instead, you must click Guide Mode *before* selecting an option from the menu bar, tool bar, or popup menu. If you change the interaction mode after selecting a configuration option, the mode change does not take effect until you select another configuration option.

Note: Guide mode is not available if your switch access level is read-only. For more information about the read-only access mode, see the “Access Modes in CMS” section on page 59.

Expert Mode

Expert mode is for users who prefer to display all the parameter fields of a feature in a single CMS window. Information about the parameter fields is available by clicking the Help button.

Wizards

Wizards simplify some configuration tasks on the switch. Similar to the guide mode, wizards provide a step-by-step approach for completing a specific configuration task. Unlike guide mode, a wizard does not prompt you to provide information for all of the

feature options. Instead, it prompts you to provide minimal information and then uses the default settings of the remaining options to set up default configurations.

Wizards are not available for all features. A menu-bar option that has wizard means that selecting that option launches the wizard for that feature.

Note: Wizards are not available if your switch access level is read-only. For more information about the read-only access mode, see the “Access Modes in CMS” section on page 59.

Tool Tips

CMS displays a popup message when you move your mouse over these devices:

- A yellow device icon in the cluster tree or in Topology view—A popup displays a fault message, such as that the RPS is faulty or that the switch is unavailable because you are in read-only mode.
- A red device icon in the cluster tree or in Topology view—A popup displays a message that the switch is down.


If you move your mouse over a table column heading, a popup displays the full heading.

Online Help

CMS provides comprehensive online help to assist you in understanding and performing configuration and monitoring tasks from the CMS windows, as shown in Figure 14. Online help includes these features:

- Feature help, available from the menu bar by selecting **Help > Contents**, provides background information and concepts on the features.
- Dialog-specific help, available from Help on the CMS windows, provides procedures for performing tasks.
- Index of help topics.
- Glossary of terms used in the online help.

Figure 14. Help Contents and Index.

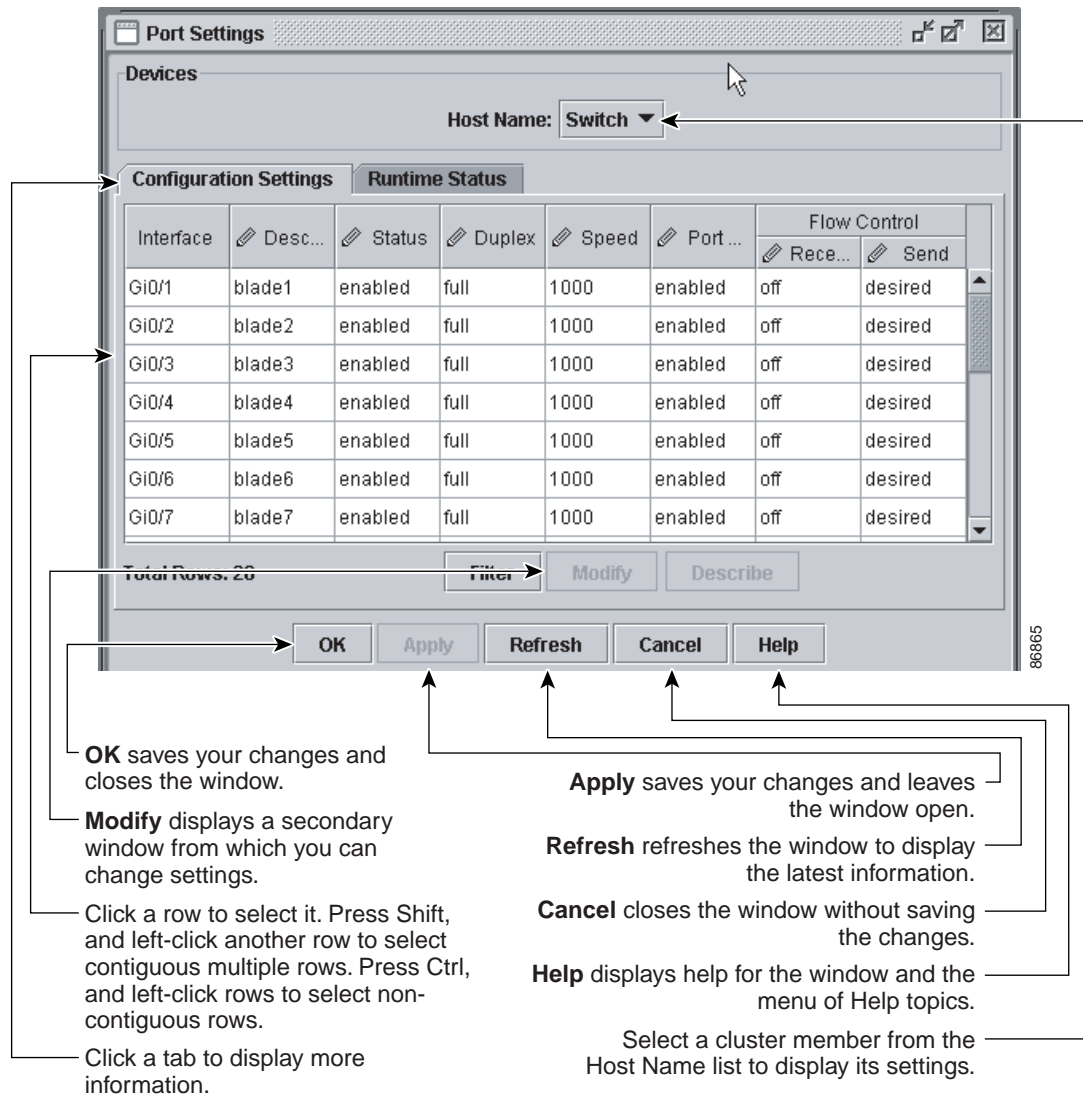


Click **Back** and **Forward** to redisplay previously displayed pages. Click **Glossary** to access the glossary from the button bar. Click **Feedback** (not shown) to send us your comments about the online help.

CMS Window Components

CMS windows consistently present configuration information. Figure 15 shows the components of a typical CMS window.

Figure 15. CMS Window Components



OK saves your changes and closes the window.

Modify displays a secondary window from which you can change settings.

Click a row to select it. Press Shift, and left-click another row to select contiguous multiple rows. Press Ctrl, and left-click rows to select non-contiguous rows.

Click a tab to display more information.

Apply saves your changes and leaves the window open.

Refresh refreshes the window to display the latest information.

Cancel closes the window without saving the changes.

Help displays help for the window and the menu of Help topics.

Select a cluster member from the Host Name list to display its settings.

Host Name List

To display or change the configuration of a cluster member, you need to select the specific switch from the Host Name drop-down list. The list appears in the configuration window of each feature and lists only the cluster members that support that feature.

Tabs, Lists, and Tables

Some CMS windows have tabs that present different sets of information. Tabs are arranged like folder headings across the top of the window. Click the tab to display its information.

Listed information can often be changed by selecting an item from a list. To change the information, select one or more items, and click **Modify**. Changing multiple items is limited to those items that apply to at least one of the selections.

Some CMS windows present information in a table format. You can edit the information in these tables.

Note: You can resize the width of the columns to display the column headings, or you can hold your cursor over the heading to display a popup description of the column.

Filter Editor

When you click **Filter** in a CMS window that contains a table, the Filter Editor window appears. The column names in the table become the field names in this window. You can enter selection criteria in these field names to filter out table rows that you do not want displayed. For procedures on using the Filter Editor, refer to the online help.

Buttons

These are the most common buttons that you use to change the information in a CMS window:

- **OK**—Save any changes and close the window. If you made no changes, the window closes. If CMS detects errors in your entry, the window remains open. For more information about error detection, see the “Red Border around a Field” section on page 58.
- **Apply**—Save any changes made in the window and leave the window open. If you made no changes, the Apply button is disabled.
- **Refresh**—Update the CMS window with the latest status of the device. Unsaved changes are lost.
- **Cancel**—Do not save any changes made in the window and close the window.
- **Help**—Display procedures on performing tasks from the window.
- **Modify**—Display the secondary window for changing information on the selected item or items. You usually select an item from a list or table and click **Modify**.

Green Border around a Field or Cell

A green border around a field or table cell means that you made an unsaved change to the field or table cell. Previous information in that field or table cell is displayed in the window status bar. When you save the changes or if you cancel the change, the green border disappears.

Red Border around a Field

A red border around a field means that you entered invalid data in the field. An error message also displays in the window status bar. When you enter valid data in the field, a green border replaces the red border until you either save or cancel the change.

If there is an error in communicating with the switch or if you make an error while performing an action, a message notifies you about the error.

Accessing CMS

This section assumes the following:

- You know the IP address and password of the command switch or a specific switch. This information is
 - Assigned to the switch through the BladeCenter Management Module WEB page, as described in the IBM BladeCenter QuickStart Guide.
 - Changed on the switch by following the information in the “Assigning Switch Information” section on page 63 and “Preventing Unauthorized Access to Your

Switch” section on page 117. Considerations for assigning IP addresses and passwords to a command switch and cluster members are described in the “IP Addresses” section on page 82 and the “Passwords” section on page 83.

- You know your access privilege level to the switch (see the “Access Modes in CMS” section on page 59).
- You have referred to the release notes for system requirements and have followed the procedures for installing the required Java plug-in and configuring your browser.

Caution: Copies of the CMS pages that you display are saved in your browser memory cache until you exit the browser session. A password is not required to redisplay these pages. You can access the command-line interface (CLI) by clicking Monitor the router - HTML access to the command line interface from a cached copy of the web page. To prevent unauthorized access to CMS and the CLI, exit your browser to end the browser session.

Note: If you have configured the Terminal Access Controller Access Control System Plus (TACACS+) or Remote Authentication Dial-In User Service (RADIUS) feature on the switch, you can still access the switch through CMS. For information about how inconsistent authentication configurations in switch clusters can affect access through CMS, see the “TACACS+ and RADIUS” section on page 83.

To access CMS, follow these steps:

1. Enter the switch IP address and your privilege level in the browser Location field (Netscape Communicator) or Address field (Microsoft Internet Explorer). For example:

`http://10.1.126.45:184/level/14/`

where 10.1.126.45 is the switch IP address, 184 is the HTTP port, and level/14 is the privilege level. You do not need to enter the HTTP port if the switch is using HTTP port 80 (the default) or enter the privilege level if you have read-write access to the switch (privilege level is 15). For information about the HTTP port, see the “HTTP Access to CMS” section on page 60. For information about privilege levels, see the “Access Modes in CMS” section on page 59.

2. When prompted for a username and password, enter only the switch enable password. CMS prompts you a second time for a username and password. Enter only the enable password again.

If you configure a local username and password, make sure you enable it by using the **ip http authentication {enable | local | tacacs}** global configuration command. Enter your username and password when prompted.

3. Click **Web Console**.

If you access CMS from a standalone or member switch, Device Manager appears. If you access CMS from a command switch, you can display the Front Panel and Topology views.

Access Modes in CMS

CMS provides two levels of access to the configuration options: read-write access and read-only access. Privilege levels 0 to 15 are supported.

- Privilege level 15 provides you with read-write access to CMS.
- Privilege levels 1 to 14 provide you with read-only access to CMS. Any options in the CMS windows, menu bar, toolbar, and popup menus that change the switch or cluster configuration are not shown in read-only mode.

- Privilege level 0 denies access to CMS.

If you do not include a privilege level when you access CMS, the switch verifies if you have privilege-level 15. If you do not, you are denied access to CMS. If you do have privilege-level 15, you are granted read-write access. Therefore, you do not need to include the privilege level if it is 15. Entering zero denies access to CMS. For more information about privilege levels, see the “Preventing Unauthorized Access to Your Switch” section on page 117.

If your cluster has these member switches running earlier software releases and if you have read-only access to these member switches, some configuration windows for those switches display incomplete information:

- Catalyst 2900 XL or Catalyst 3500 XL member switches running Cisco IOS Release 12.0(5)WC2 or earlier
- Catalyst 2950 member switches running Cisco IOS Release 12.0(5)WC2 or earlier

For more information about this limitation, refer to the release notes.

HTTP Access to CMS

CMS uses Hypertext Transfer Protocol (HTTP), which is an in-band form of communication with the switch through any one of its Ethernet ports and that allows switch management from a standard web browser. The default HTTP port is 80.

If you change the HTTP port, you must include the new port number when you enter the IP address in the browser **Location** or **Address** field (for example, `http://10.1.126.45:184` where 184 is the new HTTP port number).

Do not disable or otherwise misconfigure the port through which your management station is communicating with the switch. You might want to write down the port number to which you are connected. Changes to the switch IP information should be done with care.

For information about connecting to a switch port, refer to the switch hardware installation guide.

Saving Your Configuration

Note: The Save Configuration option is not available if your switch access level is read-only. For more information about the read-only access mode, see the “Access Modes in CMS” section on page 59.

Important: As you make cluster configuration changes (except for changes to the Topology view and in the Preferences window), make sure that you periodically save the configuration from the command switch. The configuration is saved on the command and member switches.

The front-panel images and CMS windows always display the running configuration of the switch. When you make a configuration change to a switch or switch cluster, the change becomes part of the running configuration. The change does not automatically become part of the configuration file in the flash memory, which is the startup configuration used each time the switch restarts. If you do not save your changes to the flash memory, they are lost when the switch restarts.

Restoring Your Configuration

After you save a switch configuration, you can restore the configuration to one or more switches for these reasons:

- You made an incorrect change to the current running configuration and want to reload a saved configuration.
- You need to reload a switch after a switch failure or power failure.
- You want to copy the configuration of a switch to other switches.

For CMS procedures for restoring a switch configuration, refer to the online help.

CMS Preferences

When you exit from CMS, your CMS preferences are saved to your PC in a file called `.cms_properties`. You can copy this file to other PCs. The file is stored in a default configuration directory, such as `C:\Documents and Settings\username`. If you cannot locate the CMS preferences file, select **Start > Search > For Files or Folders...**, and search for `.cms_properties`.

Note: In previous CMS versions, the preferences were saved in flash memory when you exited from CMS.

Using Different Versions of CMS

When managing switch clusters through CMS, remember that clusters can have a mix of switch models using different Cisco IOS releases and that CMS in earlier releases and on different switch platforms might look and function differently from CMS in this release.

When you select **Device > Device Manager** for a cluster member, a new browser session is launched, and the CMS version for that switch appears.

For descriptions of the CMS version that you are using, refer to the switch documentation for that specific Cisco IOS release.

Where to Go Next

Before configuring the switch, refer to these places for start-up information:

- *Cisco Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter Installation Guide*
- *Release notes for the Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter, Cisco IOS Release 12.1(14)AY*
 - CMS software requirements
 - Procedures for browser configuration
- *IBM BladeCenter QuickStart Guide* for starting up the switch through the BladeCenter Management Module WEB page
- Chapter 4 “Assigning the Switch IP Address and Default Gateway”
- Chapter 6 “Administering the Switch”

The rest of this guide provides information about the CLI procedures for the software features supported in this release. For CMS procedures and window descriptions, refer to the online help.

Chapter 4. Assigning the Switch IP Address and Default Gateway

This chapter describes how to create the initial switch configuration (for example, assign the switch IP address and default gateway information) by using a variety of automatic and manual methods.

Note: For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

This chapter consists of these sections:

- Understanding the Boot Process, on page 63
- Assigning Switch Information, on page 63
- Checking and Saving the Running Configuration, on page 65
- Modifying the Startup Configuration, on page 67
- Scheduling a Reload of the Software Image, on page 71

Understanding the Boot Process

Before you can assign switch information (IP address, subnet mask, default gateway, secret and Telnet passwords, and so forth), you need to install and power on the switch as described in the hardware installation guide that shipped with your switch.

The normal boot process involves the operation of the boot loader software, which performs these activities:

- Performs low-level CPU initialization. It initializes the CPU registers, which control where physical memory is mapped, its quantity, its speed, and so forth.
- Performs power-on self-test (POST) for the CPU subsystem. It tests the CPU DRAM and the portion of the flash device that makes up the flash file system.
- Initializes the flash file system on the system board.
- Loads a default operating system software image into memory and boots the switch.

The boot loader provides access to the flash file system before the operating system is loaded. Normally, the boot loader is used only to load, uncompress, and launch the operating system. After the boot loader gives the operating system control of the CPU, the boot loader is not active until the next system reset or power-on.

The boot loader also provides trap-door access into the system if the operating system has problems serious enough that it cannot be used. The trap-door mechanism provides enough access to the system so that if it is necessary, you can format the flash file system, reinstall the operating system software image by using the XMODEM Protocol, recover from a lost or forgotten password, and finally restart the operating system.

Assigning Switch Information

Use the BladeCenter Management Module WEB page to assign IP information to the switch. For more information, refer to the IBM BladeCenter QuickStart Guide.

If the switch reboots, the switch uses the IP address, subnet mask, and gateway configured in the stored-configuration file.

This section contains this configuration information:

- Default Switch Information, on page 64
- Manually Assigning IP Information, on page 64

Default Switch Information

Table 23 shows the default switch information.

Table 23. Default Switch Information

Feature	Default Setting
IP address	10.10.10.9x, where x is the slot number of the switch in the BladeCenter chassis.
Subnet mask	255.255.255.0.
Default gateway	No default gateway is defined.
Enable secret password	No password is defined.
Host name	The factory-assigned default host name is <i>Switch</i> .
Telnet username	USERID.
Telnet password	PASSWORD. (Note: The O is the number zero)
Cluster command switch functionality	Disabled.
Cluster name	No cluster name is defined.

Manually Assigning IP Information

Note: Before following the procedure in this section, refer to the guidelines about manually changing the switch IP address in the “Assigning Switch Information” section on page 63.

Beginning in privileged EXEC mode, follow these steps to manually assign IP information to multiple switched virtual interfaces (SVIs) or ports:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface vlan <i>vlan-id</i>	Enter interface configuration mode, and enter the VLAN to which the IP information is assigned. The range is 1 to 4094.
3.	ip address <i>ip-address subnet-mask</i>	Enter the IP address and subnet mask.
4.	exit	Return to global configuration mode.
5.	ip default-gateway <i>ip-address</i>	Enter the IP address of the next-hop router interface that is directly connected to the switch where a default gateway is being configured. The default gateway receives IP packets with unresolved destination IP addresses from the switch. Once the default gateway is configured, the switch has connectivity to the remote networks with which a host needs to communicate. Note: When your switch is configured to route with IP, it does not need to have a default gateway set.
6.	end	Return to privileged EXEC mode.
7.	show running-config	Verify your entries.
8.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

If you are removing the address through a Telnet session, your connection to the switch will be lost. To remove the default gateway address, use the **no ip default-gateway** global configuration command.

Note: The no ip address interface configuration command is not supported on the switch.

For information on setting the switch system name, protecting access to privileged EXEC commands, and setting time and calendar services, see Chapter 6 “Administering the Switch.”

Checking and Saving the Running Configuration

You can check the configuration settings that you entered or changes that you made by entering this privileged EXEC command:

```
Switch# show running-config
Building configuration...
Current configuration : 5277 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname Switch
!
!
username USERID privilege 15 password 0 PASSWORD
username USERID1 privilege 15 password 0 PASSWORD
ip subnet-zero
!
vtp mode transparent
!
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
!
vlan 2
 name operational
!
interface GigabitEthernet0/1
 description blade1
 switchport access vlan 2
 switchport trunk native vlan 2
 switchport trunk allowed vlan 2-4094
 switchport mode trunk
 storm-control broadcast level 99.99 99.98
 spanning-tree bpdufilter enable
!
interface GigabitEthernet0/2
 description blade2
 switchport access vlan 2
 switchport trunk native vlan 2
 switchport trunk allowed vlan 2-4094
 switchport mode trunk
 ip access-group SecWiz_Gi0_2_in_ip in
 spanning-tree bpdufilter enable
```

```

!
.
.
.
!
interface GigabitEthernet0/15
description mgmt1
switchport trunk allowed vlan 1
switchport mode trunk
ip access-group SecWiz_Gi0_1_out_ip in
spanning-tree cost 100
!
interface GigabitEthernet0/16
description mgmt2
switchport trunk allowed vlan 1
switchport mode trunk
ip access-group SecWiz_Gi0_1_out_ip in
spanning-tree cost 100
!
interface GigabitEthernet0/17
description extern1
switchport access vlan 2
switchport trunk native vlan 2
ip access-group SecWiz_Gi0_1_out_ip in
!
interface GigabitEthernet0/18
description extern2
switchport access vlan 2
switchport trunk native vlan 2
switchport mode access
ip access-group SecWiz_Gi0_1_out_ip in
!
interface GigabitEthernet0/19
description extern3
switchport access vlan 2
switchport trunk native vlan 2
switchport mode access
ip access-group SecWiz_Gi0_1_out_ip in
!
interface GigabitEthernet0/20
description extern4
switchport access vlan 2
switchport trunk native vlan 2
switchport mode access
ip access-group SecWiz_Gi0_1_out_ip in
speed 1000
!
interface Vlan1
ip address 172.20.138.185 255.255.255.240
no ip route-cache
!
ip default-gateway 172.20.138.178
ip http server
!
ip access-list extended SecWiz_Gi0_1_out_ip
ip access-list extended SecWiz_Gi0_2_in_ip
deny ip any host 1.1.1.1
permit ip any any
!

```

```
snmp-server community public RO
snmp-server community private RW
!
line con 0
  login local
line vty 0 4
  login local
line vty 5 15
  login local
!
end
```

To store the configuration or changes that you made to your startup configuration in flash memory, enter this privileged EXEC command:

```
Switch# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
```

This command saves the configuration settings that you made. If you fail to do this, your configuration will be lost the next time you reload the system. To display information stored in the NVRAM section of flash memory, use the **show startup-config** or **more startup-config** privileged EXEC command.

Modifying the Startup Configuration

This section describes how to modify the switch startup configuration only on the switch. It contains this configuration information:

- Default Boot Configuration, on page 67
- Specifying the Filename to Read and Write the System Configuration, on page 68
- Booting Manually, on page 68
- Booting a Specific Software Image, on page 69
- Controlling Environment Variables, on page 69

Default Boot Configuration

Table 24 shows the default boot configuration.

Table 24. Default Boot Configuration

Feature	Default Setting
Operating system software image	<p>The switch attempts to automatically boot the system using information in the BOOT environment variable. If the variable is not set, the switch attempts to load and execute the first executable image it can by performing a recursive, depth-first search throughout the flash file system.</p> <p>The software image is stored in a directory that has the same name as the image file (excluding the .bin extension).</p> <p>In a depth-first search of a directory, each encountered subdirectory is completely searched before continuing the search in the original directory.</p>
Configuration file	<p>Configured switches use the <i>config.text</i> file stored on the system board in flash memory.</p> <p>A new switch has no configuration file.</p>

Specifying the Filename to Read and Write the System Configuration

By default, the software uses the file *config.text* to read and write a nonvolatile copy of the system configuration. However, you can specify a different filename, which will be loaded during the next boot cycle.

Beginning in privileged EXEC mode, follow these steps to specify a different configuration filename:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	boot config-file flash:/file-url	<p>Specify the configuration file to load during the next boot cycle.</p> <p>For <i>file-url</i>, specify the path (directory) and the configuration filename.</p> <p>Filenames and directory names are case sensitive.</p>
3.	end	Return to privileged EXEC mode.
4.	show boot	<p>Verify your entries.</p> <p>The boot config-file global configuration command changes the setting of the CONFIG_FILE environment variable.</p>
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no boot config-file** global configuration command.

Booting Manually

By default, the switch automatically boots. This feature is not available on the IGESM.

Booting a Specific Software Image

By default, the switch attempts to automatically boot the system using information in the BOOT environment variable. If this variable is not set, the switch attempts to load and execute the first executable image it can by performing a recursive, depth-first search throughout the flash file system. In a depth-first search of a directory, each encountered subdirectory is completely searched before continuing the search in the original directory. However, you can specify a specific image to boot.

Beginning in privileged EXEC mode, follow these steps to configure the switch to boot a specific image during the next boot cycle:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	boot system <i>filesystem:/file-url</i>	Configure the switch to boot a specific image in flash memory during the next boot cycle. <ul style="list-style-type: none">For <i>filesystem:</i>, use flash: for the system board flash device.For <i>file-url</i>, specify the path (directory) and the name of the bootable image. Filenames and directory names are case sensitive.
3.	end	Return to privileged EXEC mode.
4.	show boot	Verify your entries. The boot system global command changes the setting of the BOOT environment variable. During the next boot cycle, the switch attempts to automatically boot the system using information in the BOOT environment variable.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no boot system** global configuration command.

Controlling Environment Variables

With a normally operating switch, you enter the boot loader mode only through a service port connection at 9600 bps. Use the BladeCenter management application to restart the switch. When the switch restarts, send ESC sequence characters to the service port to stop the autoboot.

You should then see the boot loader *Switch:* prompt. The boot loader performs low-level CPU initialization, performs POST, and loads a default operating system image into memory.

The switch boot loader software provides support for nonvolatile environment variables, which can be used to control how the boot loader, or any other software running on the system, behaves. Boot loader environment variables are similar to environment variables that can be set on UNIX or DOS systems.

Environment variables that have values are stored in the flash file system in various files as shown in Table 25.

Table 25. Environment Variables Storage Location

Environment Variable	Location (file system:filename)
BAUD, ENABLE_BREAK, CONFIG_BUFSIZE, CONFIG_FILE, MANUAL_BOOT, PS1	flash:env_vars
BOOT, BOOHLPR, HELPER, HELPER_CONFIG_FILE	flash:system_env_vars

Each line in these files contains an environment variable name and an equal sign followed by the value of the variable. A variable has no value if it is not listed in this file; it has a value if it is listed in the file even if the value is a null string. A variable that is set to a null string (for example, " ") is a variable with a value. Many environment variables are predefined and have default values.

Environment variables store two kinds of data:

- Data that controls code, which does not read the Cisco IOS configuration file. For example, the name of a boot loader helper file, which extends or patches the functionality of the boot loader can be stored as an environment variable.
- Data that controls code, which is responsible for reading the Cisco IOS configuration file. For example, the name of the Cisco IOS configuration file can be stored as an environment variable.

You can change the settings of the environment variables by accessing the boot loader or by using CLI commands. Under normal circumstances, it is not necessary to alter the setting of the environment variables.

Note: For complete syntax and usage information for the boot loader commands and environment variables, refer to the command reference for this release.

Table 26 describes the function of the most common environment variables.

Table 26. Environment Variables

Variable	Boot Loader Command	Cisco IOS Global Configuration Command
BOOT	<p>set BOOT <i>filesystem:/file-url ...</i></p> <p>A semicolon-separated list of executable files to try to load and execute when automatically booting. If the BOOT environment variable is not set, the system attempts to load and execute the first executable image it can find by using a recursive, depth-first search through the flash file system. If the BOOT variable is set but the specified images cannot be loaded, the system attempts to boot the first bootable file that it can find in the flash file system.</p>	<p>boot system <i>filesystem:/file-url</i></p> <p>Specifies the software image to load during the next boot cycle. This command changes the setting of the BOOT environment variable.</p>

Table 26. Environment Variables (continued)

Variable	Boot Loader Command	Cisco IOS Global Configuration Command
CONFIG_FILE	set CONFIG_FILE flash:/file-url Changes the filename that the software uses to read and write a nonvolatile copy of the system configuration.	boot config-file flash:/file-url Specifies the filename that the software uses to read and write a nonvolatile copy of the system configuration. This command changes the CONFIG_FILE environment variable.
CONFIG_BUFSIZE	set CONFIG_BUFSIZE size Changes the buffer size that the software uses to hold a copy of the configuration file in memory. The configuration file cannot be larger than the buffer size allocation. The range is from 4096 to 524288 bytes.	boot buffersize size Specifies the size of the file system-simulated NVRAM in flash memory. The buffer holds a copy of the configuration file in memory. This command changes the setting of the CONFIG_BUFSIZE environment variable. You must reload the switch by using the reload privileged EXEC command for this command to take effect.

Scheduling a Reload of the Software Image

You can schedule a reload of the software image to occur the switch at a later time (for example, late at night or during the weekend when the switch is used less), or you can synchronize a reload network-wide (for example, to perform a software upgrade on all switches in the network).

Note: A scheduled reload must take place within approximately 24 days.

Configuring a Scheduled Reload

To configure your switch to reload the software image at a later time, use one of these commands in privileged EXEC mode:

- **reload in [hh:]mm [text]**

This command schedules a reload of the software to take affect in the specified minutes or hours and minutes. The reload must take place within approximately 24 days. You can specify the reason for the reload in a string up to 255 characters in length.

- **reload at hh:mm [month day | day month] [text]**

This command schedules a reload of the software to take place at the specified time (using a 24-hour clock). If you specify the month and day, the reload is scheduled to take place at the specified time and date. If you do not specify the month and day, the reload takes place at the specified time on the current day (if the specified time is later than the current time) or on the next day (if the specified time is earlier than the current time). Specifying 00:00 schedules the reload for midnight.

Note: Use the **at** keyword only if the switch system clock has been set (through Network Time Protocol (NTP), the hardware calendar, or manually). The time is relative to the configured time zone on the switch. To schedule reloads across

several switches to occur simultaneously, the time on each switch must be synchronized with NTP.

The **reload** command halts the system. If the system is not set to manually boot, it reboots itself. Use the **reload** command after you save the switch configuration information to the startup configuration (**copy running-config startup-config**).

If your switch is configured for manual booting, do not reload it from a virtual terminal. This restriction prevents the switch from entering the boot loader mode and thereby taking it from the remote user's control.

If you modify your configuration file, the switch prompts you to save the configuration before reloading. During the save operation, the system requests whether you want to proceed with the save if the CONFIG_FILE environment variable points to a startup configuration file that no longer exists. If you proceed in this situation, the system enters setup mode upon reload.

This example shows how to reload the software on the switch on the current day at 7:30 p.m:

```
Switch# reload at 19:30
Reload scheduled for 19:30:00 UTC Wed Jun 5 1996 (in 2 hours and 25 minutes)
Proceed with reload? [confirm]
```

This example shows how to reload the software on the switch at a future time:

```
Switch# reload at 02:00 jun 20
Reload scheduled for 02:00:00 UTC Thu Jun 20 1996 (in 344 hours and 53
minutes)
Proceed with reload? [confirm]
```

To cancel a previously scheduled reload, use the **reload cancel** privileged EXEC command.

Displaying Scheduled Reload Information

To display information about a previously scheduled reload or to determine if a reload has been scheduled on the switch, use the **show reload** privileged EXEC command.

It displays reload information including the time the reload is scheduled to occur and the reason for the reload (if it was specified when the reload was scheduled).

Chapter 5. Clustering Switches

This chapter provides these topics to help you get started with switch clustering:

- Understanding Switch Clusters, on page 73
- Planning a Switch Cluster, on page 75
- Creating a Switch Cluster, on page 85
- Using the CLI to Manage Switch Clusters, on page 90
- Using SNMP to Manage Switch Clusters, on page 90

Configuring switch clusters is more easily done from the Cluster Management Suite (CMS) web-based interface than through the command-line interface (CLI). Therefore, information in this chapter focuses on using CMS to create a cluster. See Chapter 3 “Getting Started with CMS,” for additional information about switch clusters and the clustering options. For complete procedures about using CMS to configure switch clusters, refer to the online help.

For the CLI cluster commands, refer to the switch command reference.

Refer to the release notes for the list of switches eligible for switch clustering, including which ones can be command switches and which ones can only be member switches, for the required software versions, and for the browser and Java plug-in configurations.

Note: This chapter focuses on Cisco Systems Intelligent Gigabit Ethernet Switch Module switch clusters. It also includes guidelines and limitations for clusters mixed with other cluster-capable switches, but it does not provide complete descriptions of the cluster features for these other switches. For complete cluster information for a specific platform, refer to the software configuration guide for that switch.

Understanding Switch Clusters

A switch cluster is a group of connected switches that are managed as a single entity. In a switch cluster, 1 switch must be the *command switch* and up to 15 switches can be *member switches*. The total number of switches in a cluster cannot exceed 16 switches. The command switch is the single point of access used to configure, manage, and monitor the member switches. Cluster members can belong to only one cluster at a time.

When the CIGESM is used in a switch cluster, it must be managed using in-band communication at the external ports. It cannot be managed through the BladeCenter management module.

The benefits of clustering switches include:

- Management of switches regardless of their interconnection media and their physical locations. The switches can be in the same location, or they can be distributed across a Layer 2 or Layer 3 network.

Cluster members are connected to the command switch according to the connectivity guidelines described in the “Automatic Discovery of Cluster Candidates and Members” section on page 75.
- Command-switch redundancy if a command switch fails. One or more switches can be designated as *standby command switches* to avoid loss of contact with cluster members. A *cluster standby group* is a group of standby command switches.

- Management of a variety of switches through a single IP address. This conserves on IP addresses, especially if you have a limited number of them. All communication with the switch cluster is through the command switch IP address.

For other clustering benefits, see the “Advantages of Using CMS and Clustering Switches” section on page 21.

Refer to the release notes for the list of switches eligible for switch clustering, including which ones can be command switches and which ones can only be member switches, and the required software versions.

These sections describe:

- Command Switch Characteristics, on page 74
- Standby Command Switch Characteristics, on page 74
- Candidate Switch and Member Switch Characteristics, on page 74

Command Switch Characteristics

A command switch must meet these requirements:

It is running Cisco IOS Release 12.1(14)AY or later.

- It has an IP address.
- It has Cisco Discovery Protocol (CDP) version 2 enabled (the default).
- It is not a command or member switch of another cluster.

Note: The CMP-NAT-ACL access list is created when a device is configured as the command switch. Configuring any other access list on the switch can restrict access to it and affect the discovery of member and candidate switches.

Note: It is important that the CIGESM be the cluster command switch regardless of the presence of other cluster command switches in the cluster.

Standby Command Switch Characteristics

A standby command switch must meet these requirements:

- It is running Cisco IOS Release 12.1(14)AY or later.
- It has an IP address.
- It has CDP version 2 enabled.
- It is redundantly connected to the cluster so that connectivity to member switches is maintained.
- It is not a command or member switch of another cluster.

Note: When the command switch is a Cisco Systems Intelligent Gigabit Ethernet Switch Module, all standby command switches must be Cisco Systems Intelligent Gigabit Ethernet Switch Modules.

Candidate Switch and Member Switch Characteristics

Candidate switches are cluster-capable switches that have not yet been added to a cluster. *Member switches* are switches that have actually been added to a switch cluster. Although not required, a candidate or member switch can have its own IP address and password (for related considerations, see the “IP Addresses” section on page 82 and “Passwords” section on page 83).

To join a cluster, a candidate switch must meet these requirements:

- It is running cluster-capable software.
- It has CDP version 2 enabled.
- It is not a command or member switch of another cluster.

Planning a Switch Cluster

Anticipating conflicts and compatibility issues is a high priority when you manage several switches through a cluster. This section describes these guidelines, requirements, and caveats that you should understand before you create the cluster:

- Automatic Discovery of Cluster Candidates and Members, on page 75
- HSRP and Standby Command Switches, on page 79
- IP Addresses, on page 82
- Host Names, on page 82
- Passwords, on page 83
- SNMP Community Strings, on page 83
- TACACS+ and RADIUS, on page 83
- Access Modes in CMS, on page 83
- Management VLAN, on page 84
- Availability of Switch-Specific Features in Switch Clusters, on page 85

Refer to the release notes for the list of switches eligible for switch clustering, including which ones can only be member switches, for the required software versions, and for the browser and Java plug-in configurations.

Automatic Discovery of Cluster Candidates and Members

The command switch uses Cisco Discovery Protocol (CDP) to discover member switches, candidate switches, neighboring switch clusters, and edge devices in star or cascaded topologies.

Note: Do not disable CDP on the command switch, on cluster members, or on any cluster-capable switches that you might want a command switch to discover. For more information about CDP, see

Following these connectivity guidelines ensures automatic discovery of the switch cluster, cluster candidates, connected switch clusters, and neighboring edge devices:

- Discovery through CDP Hops, on page 75
- Discovery through Non-CDP-Capable and Noncluster-Capable Devices, on page 76
- Discovery through the Same Management VLAN, on page 77
- Discovery through Different Management VLANs, on page 78
- Discovery of Newly Installed Switches, on page 78

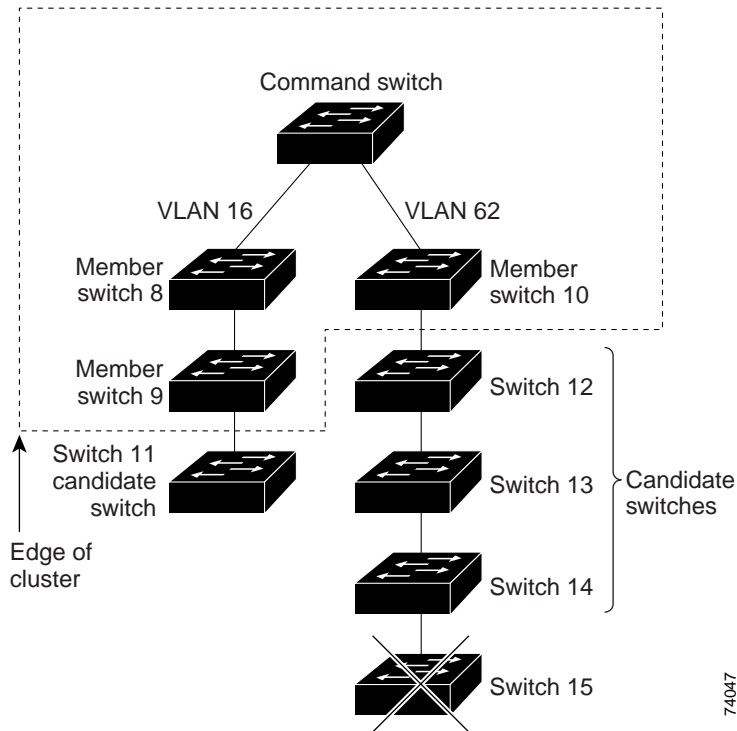
Discovery through CDP Hops

By using CDP, a command switch can discover physically connected switches up to seven CDP hops away (the default is three hops) from the edge of the cluster. The edge of the cluster is where the last member switches are connected to the cluster and to candidate switches. For example, member switches 9 and 10 in Figure 16 are at the edge of the cluster.

You can set the number of hops the command switch searches for candidate and member switches by selecting **Cluster > Hop Count**. When new candidate switches are added to the network, the command switch discovers them and adds them to the list of candidate switches.

In Figure 16, the switch has ports assigned to management VLAN 16. The CDP hop count is three. Each command switch discovers switches 11, 12, 13, and 14 because they are within three hops from the edge of the cluster. It does not discover switch 15 because it is four hops from the edge of the cluster.

Figure 16. Discovery through CDP Hops




Discovery through Non-CDP-Capable and Noncluster-Capable Devices

If a command switch is connected to a *non-CDP-capable third-party hub* (such as a non-Cisco hub), it can discover cluster-enabled devices connected to that third-party hub. However, if the command switch is connected to a *noncluster-capable Cisco device*, it cannot discover a cluster-enabled device connected beyond the noncluster-capable Cisco device.

Figure 17 shows that the command switch discovers the Catalyst 3500 XL switch, which is connected to a third-party hub. However, the command switch does not discover the IGESM that is connected to a Catalyst 5000 switch.

Figure 17. Discovery through Non-CDP-Capable and Noncluster-Capable Devices



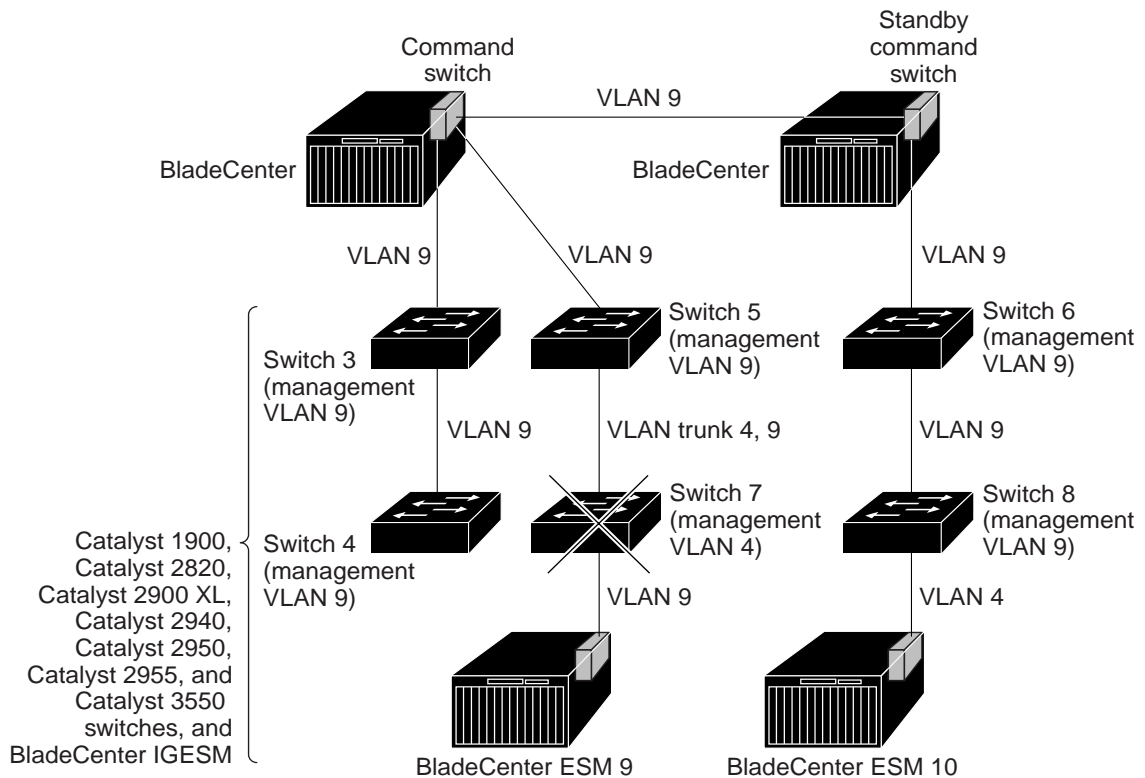
Discovery through the Same Management VLAN

A CIGESM command switch, must connect to all cluster members through its management VLAN. The default management VLAN is VLAN 1. For more information about management VLANs, see the “Management VLAN” section on page 84.

The CIGESM command switch in Figure 18 has ports assigned to management VLAN 9. It discovers all but these switches:

- Switches 7 and ESM 10 because their management VLAN (VLAN 4) is different from the command-switch management VLAN (VLAN 9)
- ESM 9 because automatic discovery does not extend beyond a noncandidate device, which is switch 7

Figure 18. Discovery through the Same Management VLAN



Discovery through Different Management VLANs

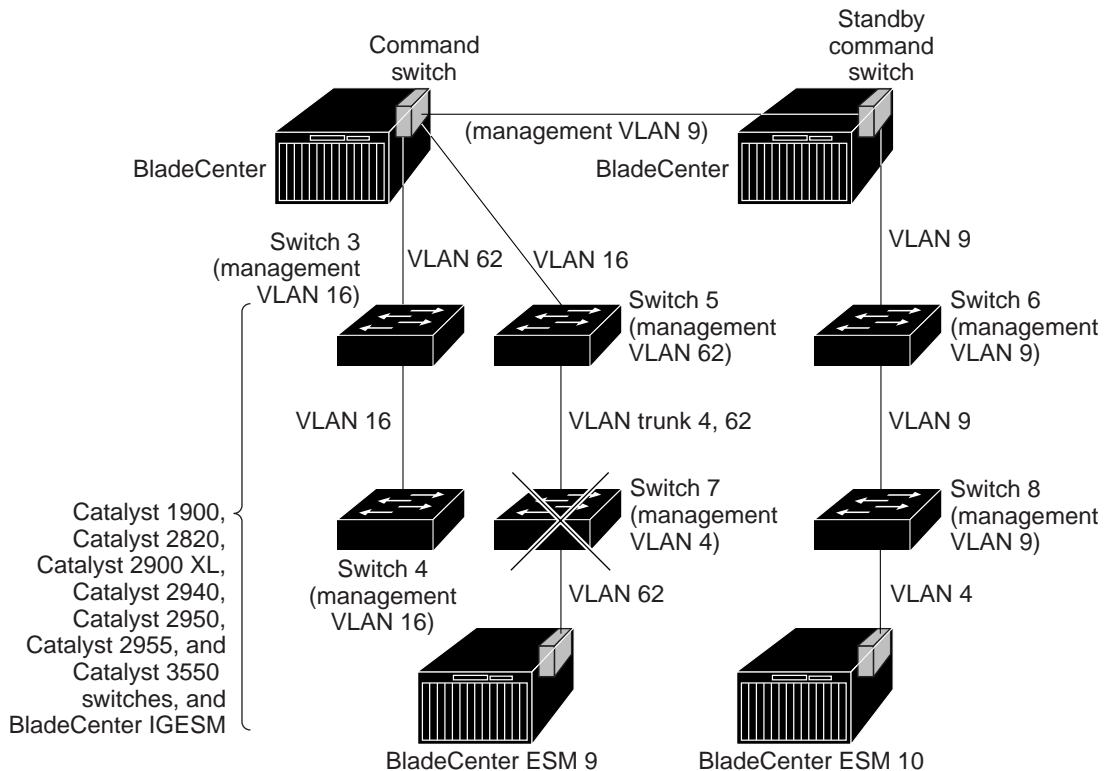
Member switches running Cisco IOS Release 12.1 (11) EA1 or later must be connected through at least one VLAN in common with the command switch. All other member switches must be connected to the command switch through their management VLAN.

In contrast, a switch running a release earlier than Cisco IOS Release (for example 295x switches) must connect to all cluster members through its management VLAN. The default management VLAN is VLAN 1. For information about discovery through the same management VLAN on these switches, see the “Discovery through the Same Management VLAN” section on page 77.

The CIGESM command switch running Cisco IOS Release 12.1(14)AY or later in Figure 19 has ports assigned to VLANs 9, 16, and 62. The management VLAN on the command switch is VLAN 9. Each command switch discovers the switches in the different management VLANs except these:

- Switches 7 and 10 (switches in management VLAN 4) because they are not connected through a common VLAN (meaning VLANs 62 and 9) with the command switch
- Switch 9 because automatic discovery does not extend beyond a noncandidate device, which is switch 7

Figure 19. Discovery through Different Management VLANs with a Layer 2 Command Switch




Discovery of Newly Installed Switches

To join a cluster, the new, out-of-the-box switch must be connected to the cluster through one of its access ports. An access port (AP) carries the traffic of and belongs to the management VLAN. By default, the new switch and its access ports are assigned to management VLAN 1.

When the new switch joins a cluster, its default management VLAN changes to the VLAN of the immediately upstream neighbor. The new switch also configures its access port to belong to the VLAN of the immediately upstream neighbor.

The command switch in Figure 20 belongs to management VLAN 16. When the switches join the cluster, their management VLAN and access ports change from VLAN 1 to VLAN 16.

Figure 20. Discovery of Newly Installed Switches in Different Management VLANs



HSRP and Standby Command Switches

The switch supports Hot Standby Router Protocol (HSRP) so that you can configure a group of standby command switches. Because a command switch manages the forwarding of all communication and configuration information to all the member switches, we strongly recommend that you configure a cluster standby command switch to take over if the primary command switch fails.

A *cluster standby group* is a group of command-capable switches that meet the requirements described in the “Standby Command Switch Characteristics” section on page 74. Only one cluster standby group can be assigned per cluster.

Note: When the command switch is a Cisco Systems Intelligent Gigabit Ethernet Switch Module, all standby command switches must be Cisco Systems Intelligent Gigabit Ethernet Switch Modules. The communication between command and standby CIGESM must be external. The Layer 2 connection will not be established through internal chassis ethernet connections.

Note: The cluster standby group is an HSRP group. Disabling HSRP disables the cluster standby group.

The switches in the cluster standby group are ranked according to HSRP priorities. The switch with the highest priority in the group is the *active command switch* (AC). The switch with the next highest priority is the *standby command switch* (SC). The other switches in the cluster standby group are the *passive command switches* (PC). If the active command switch and the standby command switch become disabled *at the same time*, the passive command switch with the highest priority becomes the active command switch. For the limitations to automatic discovery, see the “Automatic Recovery of Cluster Configuration” section on page 81. For information about

changing HSRP priority values, refer to the **standby priority** interface configuration command in the Cisco IOS Release 12.1 documentation set. The HSRP commands are the same for changing the priority of cluster standby group members and router-redundancy group members.

Note: The HSRP standby hold time interval should be greater than or equal to 3 times the hello time interval. The default HSRP standby hold time interval is 10 seconds. The default HSRP standby hello time interval is 3 seconds. For more information about the standby hold time and hello time intervals, refer to the Cisco IOS Release 12.1 documentation set on Cisco.com.

These connectivity guidelines ensure automatic discovery of the switch cluster, cluster candidates, connected switch clusters, and neighboring edge devices. These topics also provide more detail about standby command switches:

- Virtual IP Addresses, on page 80
- Other Considerations for Cluster Standby Groups, on page 80
- Automatic Recovery of Cluster Configuration, on page 81

Virtual IP Addresses

You need to assign a unique virtual IP address and group number and name to the cluster standby group. This information must be configured on the management VLAN on the active command switch. The active command switch receives traffic destined for the virtual IP address. To manage the cluster, you must access the active command switch through the virtual IP address, not through the command-switch IP address. This is in case the IP address of the active command switch is different from the virtual IP address of the cluster standby group.

If the active command switch fails, the standby command switch assumes ownership of the virtual IP address and becomes the active command switch. The passive switches in the cluster standby group compare their assigned priorities to determine the new standby command switch. The passive standby switch with the highest priority then becomes the standby command switch. When the previously active command switch becomes active again, it resumes its role as the active command switch, and the current active command switch becomes the standby command switch again. For more information about IP address in switch clusters, see the “IP Addresses” section on page 82.

Other Considerations for Cluster Standby Groups

These requirements also apply:

- When the command switch is a Cisco Systems Intelligent Gigabit Ethernet Switch Module, all standby command switches must be Cisco Systems Intelligent Gigabit Ethernet Switch Modules.
- Only one cluster standby group can be assigned to a cluster.
- All standby-group members must be members of the cluster.

Note: There is no limit to the number of switches that you can assign as standby command switches. However, the total number of switches in the cluster—which would include the active command switch, standby-group members, and member switches—cannot be more than 16.

- Each standby-group member (Figure 21) must be connected to the command switch through its management VLAN. Each standby-group member must also be redundantly connected to each other through the management VLAN.


Cisco Systems Intelligent Gigabit Ethernet Switch Modules, Catalyst 2900 XL, Catalyst 2940, Catalyst 2950, Catalyst 2955, and Catalyst 3550 member switches

must be connected to the cluster standby group through their management VLANs.

For more information about VLANs in switch clusters, see these sections:

- “Discovery through the Same Management VLAN” section on page 77
- “Discovery through Different Management VLANs” section on page 78

Figure 21. VLAN Connectivity between Standby-Group Members and Cluster Members



Automatic Recovery of Cluster Configuration

The active command switch continually forwards cluster-configuration information (but not device-configuration information) to the standby command switch. This ensures that the standby command switch can take over the cluster immediately after the active command switch fails.

Automatic discovery has these limitations:

- This limitation applies only to clusters that have Cisco Systems Intelligent Gigabit Ethernet Switch Module and Catalyst 295x command and standby command switches: If the active command switch and standby command switch become disabled *at the same time*, the passive command switch with the highest priority becomes the active command switch. However, because it was a passive standby command switch, the previous command switch *did not* forward cluster-configuration information to it. The active command switch only forwards cluster-configuration information to the standby command switch. You must therefore rebuild the cluster.
- This limitation applies to all clusters: If the active command switch fails and there are more than two switches in the cluster standby group, the new command switch does not discover any Catalyst 2916M XL member switches. You must re-add these member switches to the cluster.

- This limitation applies to all clusters: If the active command switch fails and becomes active again, it does not discover any Catalyst 2916M XL member switches. You must again add these member switches to the cluster.

When the previously active command switch resumes its active role, it receives a copy of the latest cluster configuration from the active command switch, including members that were added while it was down. The active command switch sends a copy of the cluster configuration to the cluster standby group.

IP Addresses

You must assign IP information to a command switch. You can access the cluster through the command-switch IP address. If you configure a cluster standby group, you must use the standby-group virtual IP address to manage the cluster from the active command switch. Using the virtual IP address ensures that you retain connectivity to the cluster if the active command switch fails and that a standby command switch becomes the active command switch.

If the active command switch fails and the standby command switch takes over, you must either use the standby-group virtual IP address or the IP address available on the new active command switch to access the cluster.

You can assign an IP address to a cluster-capable switch, but it is not necessary. A member switch is managed and communicates with other member switches through the command-switch IP address. If the member switch leaves the cluster and it does not have its own IP address, you then must assign IP information to it to manage it as a standalone switch.

Note: Changing the command switch IP address ends your CMS session on the switch. Restart your CMS session by entering the new IP address in the browser **Location** field (Netscape Communicator) or **Address** field (Internet Explorer), as described in the “Launching CMS” section on page 35.

For more information about IP addresses, see Chapter 4 “Assigning the Switch IP Address and Default Gateway.”

Host Names

You do not need to assign a host name to either a command switch or an eligible cluster member. However, a host name assigned to the command switch can help to identify the switch cluster. The default host name for the switch is *Switch*.

If a switch joins a cluster and it does not have a host name, the command switch appends a unique member number to its own host name and assigns it sequentially as each switch joins the cluster. The number means the order in which the switch was added to the cluster. For example, a command switch named *eng-cluster* could name the fifth cluster member *eng-cluster-5*.

If a switch has a host name, it retains that name when it joins a cluster. It retains that host name even after it leaves the cluster.

If a switch received its host name from the command switch, was removed from a cluster, was then added to a new cluster, and kept the same member number (such as 5), the old host name (such as *eng-cluster-5*) is overwritten with the host name of the command switch in the new cluster (such as *mkg-cluster-5*). If the switch member number changes in the new cluster (such as 3), the switch retains the previous name (*eng-cluster-5*).

Passwords

You do not need to assign passwords to an individual switch if it will be a cluster member. When a switch joins a cluster, it inherits the command-switch password and retains it when it leaves the cluster. If no command-switch password is configured, the member switch inherits a null password. Member switches only inherit the command-switch password.

If you change the member-switch password to be different from the command-switch password and save the change, the switch is not manageable by the command switch until you change the member-switch password to match the command-switch password. Rebooting the member switch does not revert the password back to the command-switch password. We recommend that you do not change the member-switch password after it joins a cluster.

For more information about passwords, see the “Preventing Unauthorized Access to Your Switch” section on page 117.

SNMP Community Strings

A member switch inherits the command-switch first read-only (RO) and read-write (RW) community strings with *@esN* appended to the community strings:

- *command-switch-readonly-community-string@esN*, where *N* is the member-switch number.
- *command-switch-readwrite-community-string@esN*, where *N* is the member-switch number.

If the command switch has multiple read-only or read-write community strings, only the first read-only and read-write strings are propagated to the member switch.

The switches support an unlimited number of community strings and string lengths.

TACACS+ and RADIUS

Inconsistent authentication configurations in switch clusters cause CMS to continually prompt for a user name and password. If Terminal Access Controller Access Control System Plus (TACACS+) is configured on a cluster member, it must be configured on all cluster members. Similarly, if Remote Authentication Dial-In User Service (RADIUS) is configured on a cluster member, it must be configured on all cluster members. Further, the same switch cluster cannot have some members configured with TACACS+ and other members configured with RADIUS.

For more information about TACACS+, see the “Controlling Switch Access with TACACS+” section on page 124. For more information about RADIUS, see the “Controlling Switch Access with RADIUS” section on page 131.

Access Modes in CMS

CMS provides two levels of access to the configuration options: read-write access and read-only access. Privilege levels 0 to 15 are supported.

- Privilege level 15 provides you with read-write access to CMS.
- Privilege levels 1 to 14 provide you with read-only access to CMS. Any options in the CMS windows, menu bar, toolbar, and popup menus that change the switch or cluster configuration are not shown in read-only mode.
- Privilege level 0 denies access to CMS.

For more information about CMS access modes, see the “Access Modes in CMS” section on page 59.

Note: If your cluster has these member switches running earlier software releases and if you have read-only access to these member switches, some configuration windows for those switches display incomplete information:

- Catalyst 2900 XL or Catalyst 3500 XL member switches running Cisco IOS Release 12.0(5)WC2 or earlier
- Non-LRE Catalyst 2950 member switches running Cisco IOS Release 12.0(5)WC2 or earlier
- Catalyst 3550 member switches running Cisco IOS Release 12.1(6)EA1 or earlier

For more information about this limitation, refer to the release notes.

- Catalyst 2900 XL switches with 4-MB CPU DRAM do not support read-only mode on CMS

In read-only mode, these switches appear as unavailable devices and cannot be configured from CMS.

Management VLAN

Communication with the switch management interfaces is through the command-switch IP address. The IP address is associated with the management VLAN, which by default is VLAN 1. To manage switches in a cluster, the command switch, member switches, and candidate switches must be connected through ports assigned to the command-switch management VLAN.

Note: If the command switch is a Catalyst 2950 running Cisco IOS Release 12.1(9)EA1 or later or a Catalyst 2955, candidate and member switches can belong to different management VLANs. However, they must connect to the command switch through their management VLAN.

- Catalyst 2950 standby command switches running Cisco IOS Release 12.1(9)EA1 or later and Catalyst 2955 standby command switches can connect to candidate and member switches in VLANs different from their management VLANs.

If you add a new, out-of-box switch to a cluster and the cluster is using a management VLAN other than the default VLAN 1, the command switch automatically senses that the new switch has a different management VLAN and has not been configured. The command switch issues commands to change the management VLAN of the new switch to the one the cluster is using. This automatic VLAN change only occurs for new, out-of-box switches that do not have a config.text file and that have no changes to the running configuration. For more information, see the “Discovery of Newly Installed Switches” section on page 78.

You can change the management VLAN of a member switch (not the command switch). However, the command switch will not be able to communicate with it. In this case, you will need to manage the switch as a standalone switch.

You can globally change the management VLAN for the cluster as long as each member switch has either a trunk connection or a connection to the new command-switch management VLAN. From the command switch, use the **cluster management vlan** global configuration command to change the cluster management VLAN to a different management VLAN.

Caution: You can change the management VLAN through a service port connection without interrupting the connection. However, changing the management VLAN ends your CMS session. Restart your CMS session by

entering the new IP address in the browser Location field (Netscape Communicator) or Address field (Microsoft Internet Explorer), as described in the “HTTP Access to CMS” section on page 60.

For more information about changing the management VLAN, see the “Management VLAN” section on page 84.

Availability of Switch-Specific Features in Switch Clusters

The menu bar on the command switch displays all options available from the switch cluster. Therefore, features specific to a member switch are available from the command-switch menu bar. For example, **Device > LRE Profile** appears in the command-switch menu bar when at least one Catalyst 2900 LRE XL or Catalyst 2950 LRE switch is in the cluster.

Creating a Switch Cluster

Using CMS to create a cluster is easier than using the CLI commands. This section provides this information:

- Enabling a Command Switch, on page 85
- Adding Member Switches, on page 86
- Creating a Cluster Standby Group, on page 87
- Verifying a Switch Cluster, on page 89

This section assumes you have already cabled the switches, as described in the switch hardware installation guide, and followed the guidelines described in the “Planning a Switch Cluster” section on page 75.


Note: Refer to the release notes for the list of Catalyst switches eligible for switch clustering, including which ones can be command switches and which ones can only be member switches, for the required software versions, and for the browser and Java plug-in configurations.

Enabling a Command Switch

The switch you designate as the command switch must meet the requirements described in the “Command Switch Characteristics” section on page 74, the “Planning a Switch Cluster” section on page 75, and the release notes.

If you did not enable a command switch during initial switch setup, launch Device Manager from a command-capable switch, and select **Cluster > Create Cluster**. Enter a cluster number (the default is 0), and use up to 31 characters to name the cluster (Figure 22). Instead of using CMS to enable a command switch, you can use the **cluster enable** global configuration command.

Figure 22. Create Cluster Window



Adding Member Switches

As explained in the “Automatic Discovery of Cluster Candidates and Members” section on page 75, the command switch automatically discovers candidate switches. When you add new cluster-capable switches to the network, the command switch discovers them and adds them to a list of candidate switches. To display an updated cluster candidates list from the Add to Cluster window (Figure 23), either relaunch CMS and redisplay this window, or follow these steps:

1. Close the Add to Cluster window.
2. Select **View > Refresh**.
3. Select **Cluster > Add to Cluster** to redisplay the Add to Cluster window.

From CMS, there are two ways to add switches to a cluster:

- Select **Cluster > Add to Cluster**, select a candidate switch from the list, click **Add**, and click **OK**. To add more than one candidate switch, press **Ctrl**, and make your choices, or press **Shift**, and choose the first and last switch in a range.
- Display the Topology view, right-click a candidate-switch icon, and select **Add to Cluster** (Figure 24). In the Topology view, candidate switches are cyan, and member switches are green. To add more than one candidate switch, press **Ctrl**, and left-click the candidates that you want to add.

Instead of using CMS to add members to the cluster, you can use the **cluster member** global configuration command from the command switch. Use the **password** option in this command if the candidate switch has a password.

You can select 1 or more switches as long as the total number of switches in the cluster does not exceed 16 (this includes the command switch). When a cluster has 16 members, the **Add to Cluster** option is not available for that cluster. In this case, you must remove a member switch before adding a new one.

If a password has been configured on a candidate switch, you are prompted to enter it before it can be added to the cluster. If the candidate switch does not have a password, any entry is ignored.

If multiple candidate switches have the same password, you can select them as a group, and add them at the same time.

If a candidate switch in the group has a password different from the group, only that specific candidate switch is not added to the cluster.

When a candidate switch joins a cluster, it inherits the command-switch password. For more information about setting passwords, see the “Passwords” section on page 83.

For additional authentication considerations in switch clusters, see the “TACACS+ and RADIUS” section on page 83.

Figure 23. Add to Cluster Window

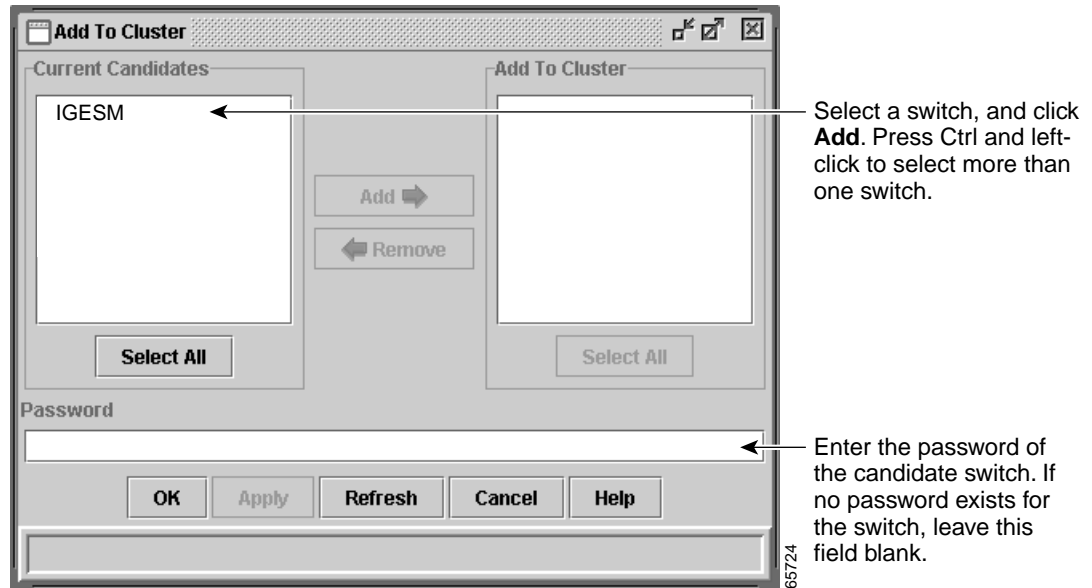
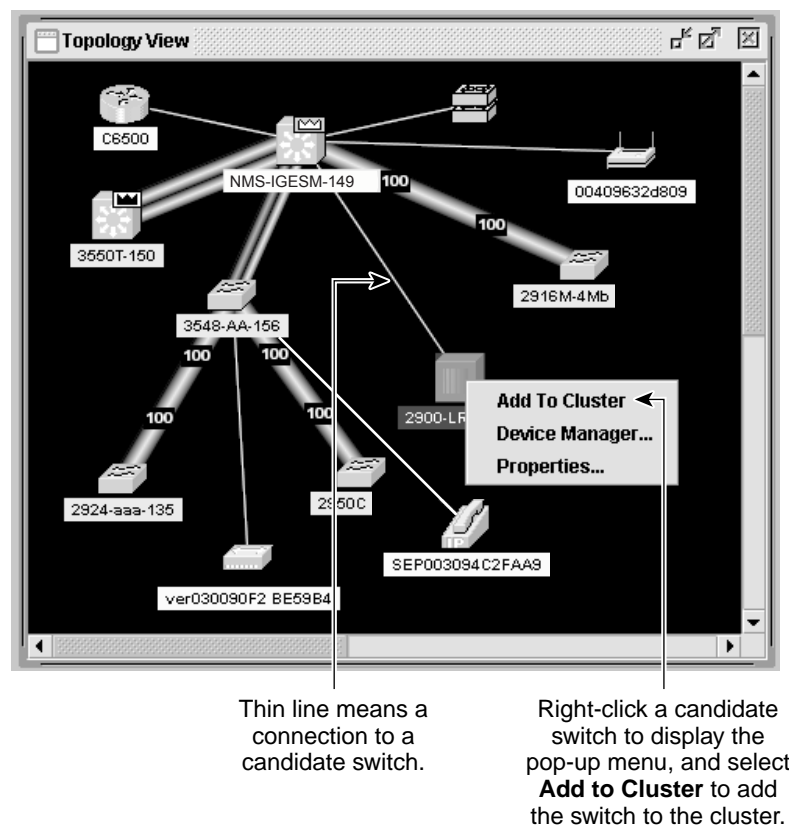


Figure 24. Using the Topology View to Add Member Switches



Creating a Cluster Standby Group

The cluster standby group members must meet the requirements described in the “Standby Command Switch Characteristics” section on page 74 and “HSRP and

Standby Command Switches” section on page 79. To create a cluster standby group, select **Cluster > Standby Command Switches** (Figure 25).

Instead of using CMS to add switches to a standby group and to bind the standby group to a cluster, you can use the **standby ip**, the **standby name**, and the **standby priority** interface configuration commands and the **cluster standby group** global configuration command.

Note: When the command switch is a Cisco Systems Intelligent Gigabit Ethernet Switch Module, all standby command switches must be Cisco Systems Intelligent Gigabit Ethernet Switch Modules.

- When the command switch is a Catalyst 2955 switch, all standby command switches must be Catalyst 2955 switches.
- When the command switch is a Catalyst 2950 LRE switch, all standby command switches must be Catalyst 2950 LRE switches.
- When the command switch is a non-LRE Catalyst 2950 switch running Cisco IOS Release 12.1(9)EA1 or later, all standby command switches must be non-LRE Catalyst 2950 switches running Cisco IOS Release 12.1(9)EA1 or later.
- When the command switch is a non-LRE Catalyst 2950 switch running Cisco IOS Release 12.1(6)EA2 or later, all standby command switches must be non-LRE Catalyst 2950 switches running Cisco IOS Release 12.1(6)EA2 or later.
- When the command switch is running Cisco IOS Release 12.0(5)WC2 or earlier, the standby command switches can be these switches: Catalyst 2900 XL, non-LRE Catalyst 2950, and Catalyst 3500 XL switches.

These abbreviations are appended to the switch host names in the Standby Command Group list to show their eligibility or status in the cluster standby group:

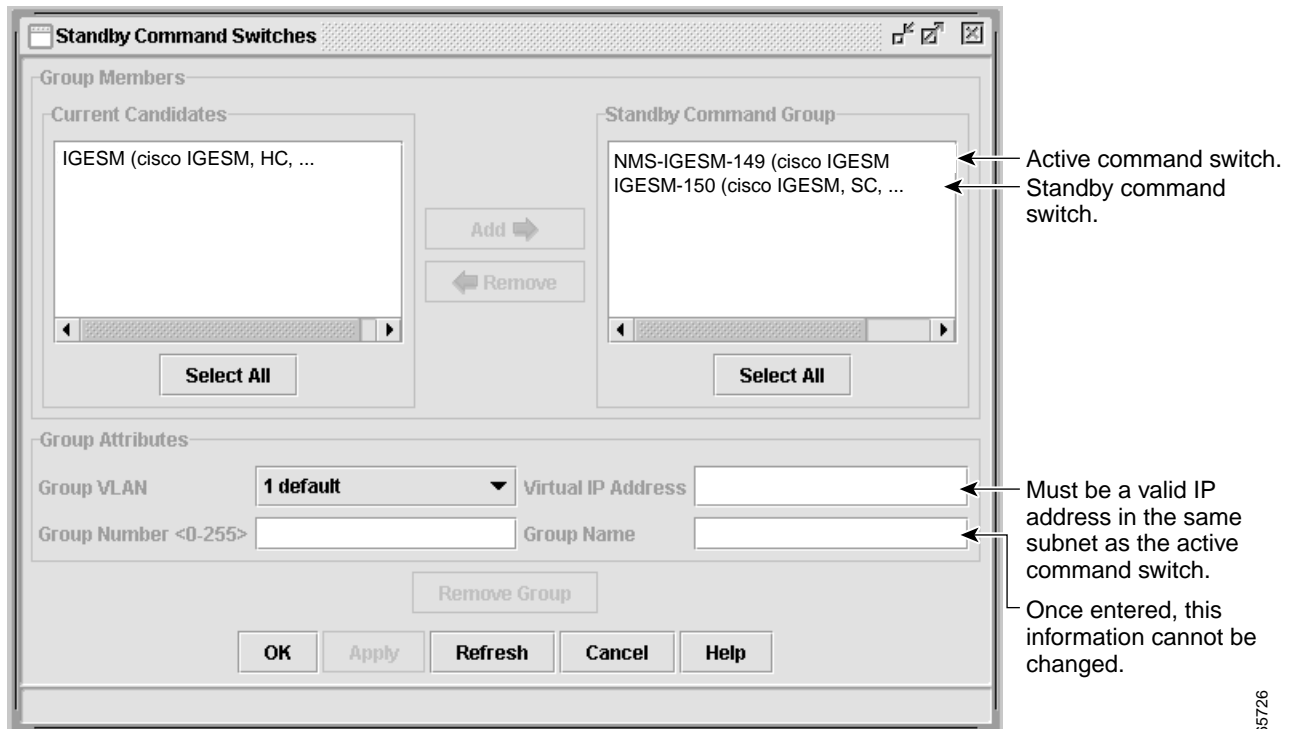
- AC—Active command switch
- SC—Standby command switch
- PC—Member of the cluster standby group but not the standby command switch
- HC—Candidate switch that can be added to the cluster standby group
- CC—Command switch when HSRP is disabled

You must enter a virtual IP address for the cluster standby group. This address must be in the same subnet as the IP addresses of the switch. The group number must be unique within the IP subnet. It can be from 0 to 255, and the default is 0. The group name can have up to 31 characters.

The Standby Command Configuration window uses the default values for the **preempt** and **name** commands that you have set by using the CLI. If you use this window to create the HSRP group, all switches in the group have the **preempt** command enabled. You must also provide a name for the group.

Note: The HSRP standby hold time interval should be greater than or equal to 3 times the hello time interval. The default HSRP standby hold time interval is 10 seconds. The default HSRP standby hello time interval is 3 seconds. For more information about the standby hold time and hello time intervals, refer to the Cisco IOS Release 12.1 documentation set on Cisco.com.

Figure 25. Standby Command Configuration Window.



Verifying a Switch Cluster

When you finish adding cluster members, follow these steps to verify the cluster:

1. Enter the command switch IP address in the browser **Location** field (Netscape Communicator) or **Address** field (Microsoft Internet Explorer) to access all switches in the cluster.
2. Enter the command-switch password.
3. Select **View > Topology** to display the cluster topology and to view link information. For complete information about the Topology view, including descriptions of the icons, links, and colors, see the “Topology View” section on page 42.
4. Select **Reports > Inventory** to display an inventory of the switches in the cluster (Figure 26).

The summary includes information such as switch model numbers, serial numbers, software versions, IP information, and location.

You can also display port and switch statistics from **Reports > Port Statistics** and **Port > Port Settings > Runtime Status**.

Instead of using CMS to verify the cluster, you can use the **show cluster members** user EXEC command from the command switch or use the **show cluster** user EXEC command from the command switch or from a member switch.

65726

Figure 26. Inventory Window

Host Name	Device Type	Serial Number	IP Address	Software Version	Sys Location	Module 1	Module 2
NMS-IGESM-149	cisco WS-IGESM		10.1.1.2, 10.10.10.5	12.1(14)AY		NA	NA
3548-AA-156	cisco WS-C3548-XL	FAA0428X096	10.10.10.6	12.0(5.3)WC(1)		NA	NA
2900-LRE-24-1	cisco WS-C2924-LRE-XL	FAA0514E08M	10.10.10.7	12.0(5)WC2		NA	NA
IGESM-150	cisco WS-IGESM	FAA0514E07W	10.1.1.2, 10.10.10.1, 10.	12.1(14)AY	SJ	NA	NA
2916M-4Mb	cisco WS-C2916M-XL	FAA0306S0NY	10.10.10.2	11.2(8.6)SA6		100BTX	WS-X2914-XL-V
2950C	cisco WS-C2950C-24	FAB0517Q0F7	10.10.10.3	12.1(6)EA2		NA	NA
2924-aaa-135	cisco WS-C2924-XL	FAB0433V0E2	10.10.10.9	13.0(5)XU		NA	NA

Buttons: OK, Refresh, Help

For more information about creating and managing clusters, refer to the online help. For information about the cluster commands, refer to the switch command reference.

Using the CLI to Manage Switch Clusters

You can configure member switches from the CLI by first logging into the command switch. Enter the **rcommand** user EXEC command and the member switch number to start a Telnet session (through a Telnet connection) and to access the member switch CLI. The command mode changes, and the CLI commands operate as usual. Enter the **exit** privileged EXEC command on the member switch to return to the command-switch CLI.

This example shows how to log into member-switch 3 from the command-switch CLI:

```
swit ch# rcommand 3
```

If you do not know the member-switch number, enter the **show cluster members** privileged EXEC command on the command switch. For more information about the **rcommand** command and all other cluster commands, refer to the switch command reference.

The Telnet session accesses the member-switch CLI at the same privilege level as on the command switch. The CLI commands then operate as usual.

Using SNMP to Manage Switch Clusters

When you create a cluster, the command switch manages the exchange of messages between member switches and an SNMP application. The cluster software on the command switch appends the member switch number (**@esN**, where **N** is the switch number) to the first configured read-write and read-only community strings on the command switch and propagates them to the member switch. The command switch uses this community string to control the forwarding of gets, sets, and get-next messages between the SNMP management station and the member switches.


Note: When a cluster standby group is configured, the command switch can change without your knowledge. Use the first read-write and read-only community strings to communicate with the command switch if there is a cluster standby group configured for the cluster.

If the member switch does not have an IP address, the command switch redirects traps from the member switch to the management station, as shown in Figure 27. If a

member switch has its own IP address and community strings, the member switch can send traps directly to the management station, without going through the command switch.

If a member switch has its own IP address and community strings, they can be used in addition to the access provided by the command switch. For more information about SNMP and community strings, see Chapter 23 "Configuring SNMP."

Figure 27. SNMP Management for a Cluster



Chapter 6. Administering the Switch

This chapter describes how to perform one-time operations to administer your switch. This chapter consists of these sections:

- Managing the System Time and Date, on page 93
- Configuring a System Name and Prompt, on page 105
- Creating a Banner, on page 108
- Managing the MAC Address Table, on page 110
- Managing the ARP Table, on page 116

Managing the System Time and Date

You can manage the system time and date on your switch using automatic configuration, such as the Network Time Protocol (NTP), or manual configuration methods.

Note: For complete syntax and usage information for the commands used in this section, refer to the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference*.

This section contains this configuration information:

- Understanding the System Clock, on page 93
- Understanding Network Time Protocol, on page 94
- Configuring NTP, on page 95
- Configuring Time and Date Manually, on page 102

Understanding the System Clock

The heart of the time service is the system clock. This clock runs from the moment the system starts up and keeps track of the date and time.

The system clock can then be set from these sources:

- Network Time Protocol
- Manual configuration

The system clock can provide time to these services:

- User **show** commands
- Logging and debugging messages

The system clock keeps track of time internally based on Universal Time Coordinated (UTC), also known as Greenwich Mean Time (GMT). You can configure information about the local time zone and summer time (daylight saving time) so that the time is correctly displayed for the local time zone.

The system clock keeps track of whether the time is *authoritative* or not (that is, whether it has been set by a time source considered to be authoritative). If it is not authoritative, the time is available only for display purposes and is not redistributed. For configuration information, see the “Configuring Time and Date Manually” section on page 102.

Understanding Network Time Protocol

The NTP is designed to time-synchronize a network of devices. NTP runs over User Datagram Protocol (UDP), which runs over IP. NTP is documented in RFC 1305.

An NTP network usually gets its time from an authoritative time source, such as a radio clock or an atomic clock attached to a time server. NTP then distributes this time across the network. NTP is extremely efficient; no more than one packet per minute is necessary to synchronize two devices to within a millisecond of one another.

NTP uses the concept of a *stratum* to describe how many NTP hops away a device is from an authoritative time source. A stratum 1 time server has a radio or atomic clock directly attached, a stratum 2 time server receives its time through NTP from a stratum 1 time server, and so on. A device running NTP automatically chooses as its time source the device with the lowest stratum number with which it communicates through NTP. This strategy effectively builds a self-organizing tree of NTP speakers.

NTP avoids synchronizing to a device whose time might not be accurate by never synchronizing to a device that is not synchronized. NTP also compares the time reported by several devices and does not synchronize to a device whose time is significantly different than the others, even if its stratum is lower.

The communications between devices running NTP (known as *associations*) are usually statically configured; each device is given the IP address of all devices with which it should form associations. Accurate timekeeping is possible by exchanging NTP messages between each pair of devices with an association. However, in a LAN environment, NTP can be configured to use IP broadcast messages instead. This alternative reduces configuration complexity because each device can simply be configured to send or receive broadcast messages. However, in that case, information flow is one-way only.

The time kept on a device is a critical resource; you should use the security features of NTP to avoid the accidental or malicious setting of an incorrect time. Two mechanisms are available: an access list-based restriction scheme and an encrypted authentication mechanism.

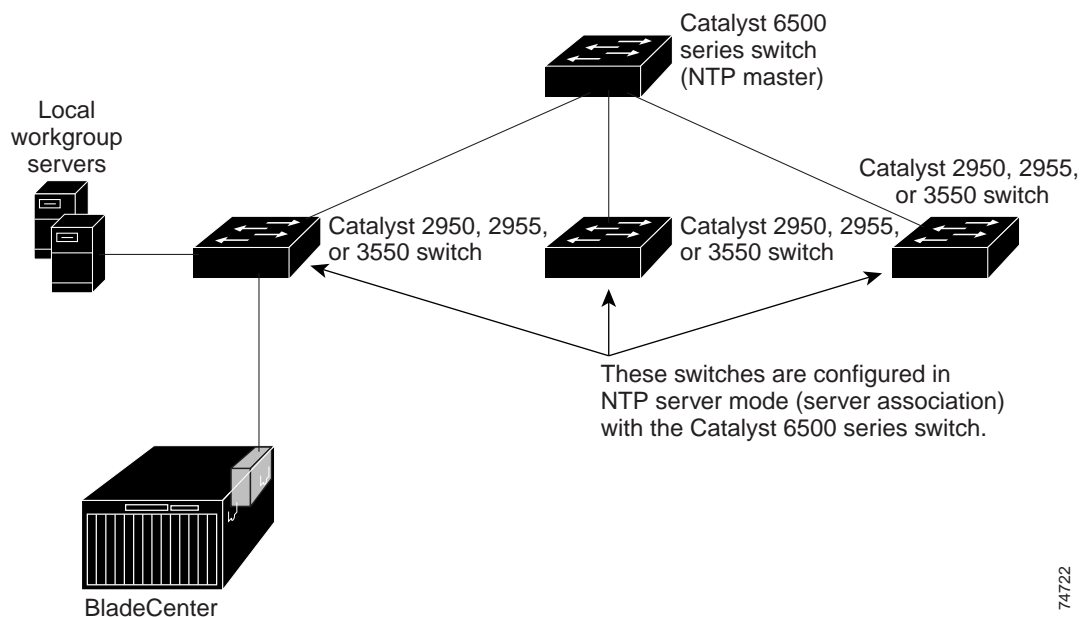
Cisco's implementation of NTP does not support stratum 1 service; it is not possible to connect to a radio or atomic clock. We recommend that the time service for your network be derived from the public NTP servers available on the IP Internet. Figure 28 shows a typical network example using NTP.

If the network is isolated from the Internet, Cisco's implementation of NTP allows a device to act as though it is synchronized through NTP, when in fact it has determined the time by using other means. Other devices then synchronize to that device through NTP.

When multiple sources of time are available, NTP is always considered to be more authoritative. NTP time overrides the time set by any other method.

Several manufacturers include NTP software for their host systems, and a publicly available version for systems running UNIX and its various derivatives is also available. This software allows host systems to be time-synchronized as well.

Figure 28. Typical NTP Network Configuration



74722

Configuring NTP

The switch does not have a hardware-supported clock, and it cannot function as an NTP master clock to which peers synchronize themselves when an external NTP source is not available. The switch also has no hardware support for a calendar. As a result, the **ntp update-calendar** and the **ntp master** global configuration commands are not available.

This section contains this configuration information:

- Default NTP Configuration, on page 95
- Configuring NTP Authentication, on page 96
- Configuring NTP Associations, on page 97
- Configuring NTP Broadcast Service, on page 98
- Configuring NTP Access Restrictions, on page 99
- Configuring the Source IP Address for NTP Packets, on page 101
- Displaying the NTP Configuration, on page 102

Default NTP Configuration

Table 27 shows the default NTP configuration.

Table 27. Default NTP Configuration

Feature	Default Setting
NTP authentication	Disabled. No authentication key is specified.
NTP peer or server associations	None configured.
NTP broadcast service	Disabled; no interface sends or receives NTP broadcast packets.
NTP access restrictions	No access control is specified.
NTP packet source IP address	The source address is determined by the outgoing interface.

NTP is enabled on all interfaces by default. All interfaces receive NTP packets.

Configuring NTP Authentication

This procedure must be coordinated with the administrator of the NTP server; the information you configure in this procedure must be matched by the servers used by the switch to synchronize its time to the NTP server.

Beginning in privileged EXEC mode, follow these steps to authenticate the associations (communications between devices running NTP that provide for accurate timekeeping) with other devices for security purposes:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	ntp authenticate	Enable the NTP authentication feature, which is disabled by default.
3.	ntp authentication-key <i>number</i> md5 <i>value</i>	Define the authentication keys. By default, none are defined. <ul style="list-style-type: none"> For <i>number</i>, specify a key number. The range is 1 to 4294967295. md5 specifies that message authentication support is provided by using the message digest algorithm 5 (MD5). For <i>value</i>, enter an arbitrary string of up to eight characters for the key. <p>The switch does not synchronize to a device unless both have one of these authentication keys, and the key number is specified by the ntp trusted-key <i>key-number</i> command.</p>
4.	ntp trusted-key <i>key-number</i>	Specify one or more key numbers (defined in Step 3) that a peer NTP device must provide in its NTP packets for this switch to synchronize to it. <p>By default, no trusted keys are defined.</p> <p>For <i>key-number</i>, specify the key defined in Step 3.</p> <p>This command provides protection against accidentally synchronizing the switch to a device that is not trusted.</p>
5.	end	Return to privileged EXEC mode.
6.	show running-config	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable NTP authentication, use the **no ntp authenticate** global configuration command. To remove an authentication key, use the **no ntp authentication-key** *number* global configuration command. To disable authentication of the identity of a device, use the **no ntp trusted-key** *key-number* global configuration command.

This example shows how to configure the switch to synchronize only to devices providing authentication key 42 in the device's NTP packets:

```
Switch(config)# ntp authenticate
Switch(config)# ntp authentication-key 42 md5 aNiceKey
Switch(config)# ntp trusted-key 42
```

Configuring NTP Associations

An NTP association can be a peer association (this switch can either synchronize to the other device or allow the other device to synchronize to it), or it can be a server association (meaning that only this switch synchronizes to the other device, and not the other way around).

Beginning in privileged EXEC mode, follow these steps to form an NTP association with another device:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	ntp peer <i>ip-address</i> [version <i>number</i>] [key <i>keyid</i>] [source <i>interface</i>] [prefer] or ntp server <i>ip-address</i> [version <i>number</i>] [key <i>keyid</i>] [source <i>interface</i>] [prefer]	Configure the switch system clock to synchronize a peer or to be synchronized by a peer (peer association). or Configure the switch system clock to be synchronized by a time server (server association). No peer or server associations are defined by default. <ul style="list-style-type: none"> For <i>ip-address</i> in a peer association, specify either the IP address of the peer providing, or being provided, the clock synchronization. For a server association, specify the IP address of the time server providing the clock synchronization. (Optional) For <i>number</i>, specify the NTP version number. The range is 1 to 3. By default, version 3 is selected. (Optional) For <i>keyid</i>, enter the authentication key defined with the ntp authentication-key global configuration command. (Optional) For <i>interface</i>, specify the interface from which to pick the IP source address. By default, the source IP address is taken from the outgoing interface. (Optional) Enter the prefer keyword to make this peer or server the preferred one that provides synchronization. This keyword reduces switching back and forth between peers and servers.
3.	end	Return to privileged EXEC mode.
4.	show running-config	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

You need to configure only one end of an association; the other device can automatically establish the association. If you are using the default NTP version (version 3) and NTP synchronization does not occur, try using NTP version 2. Many NTP servers on the Internet run version 2.

To remove a peer or server association, use the **no ntp peer ip-address** or the **no ntp server ip-address** global configuration command.

This example shows how to configure the switch to synchronize its system clock with the clock of the peer at IP address 172.16.22.44 using NTP version 2:

```
Switch(config)# ntp server 172.16.22.44 version 2
```

Configuring NTP Broadcast Service

The communications between devices running NTP (known as *associations*) are usually statically configured; each device is given the IP addresses of all devices with which it should form associations. Accurate timekeeping is possible by exchanging NTP messages between each pair of devices with an association. However, in a LAN environment, NTP can be configured to use IP broadcast messages instead. This alternative reduces configuration complexity because each device can simply be configured to send or receive broadcast messages. However, the information flow is one-way only.

The switch can send or receive NTP broadcast packets on an interface-by-interface basis if there is an NTP broadcast server, such as a router, broadcasting time information on the network. The switch can send NTP broadcast packets to a peer so that the peer can synchronize to it. The switch can also receive NTP broadcast packets to synchronize its own clock. This section has procedures for both sending and receiving NTP broadcast packets.

Beginning in privileged EXEC mode, follow these steps to configure the switch to send NTP broadcast packets to peers so that they can synchronize their clock to the switch:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Specify the interface to send NTP broadcast packets, and enter interface configuration mode.
3.	ntp broadcast [version <i>number</i>] [key <i>keyid</i>] [<i>destination-address</i>]	Enable the interface to send NTP broadcast packets to a peer. By default, this feature is disabled on all interfaces. <ul style="list-style-type: none">• (Optional) For <i>number</i>, specify the NTP version number. The range is 1 to 3. If you do not specify a version, version 3 is used.• (Optional) For <i>keyid</i>, specify the authentication key to use when sending packets to the peer.• (Optional) For <i>destination-address</i>, specify the IP address of the peer that is synchronizing its clock to this switch.
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.
7.		Configure the connected peers to receive NTP broadcast packets as described in the next procedure.

To disable the interface from sending NTP broadcast packets, use the **no ntp broadcast** interface configuration command.

This example shows how to configure an interface to send NTP version 2 packets:

```
Switch(config)# interface gigabitethernet0/17  
Switch(config-if)# ntp broadcast version 2
```

Beginning in privileged EXEC mode, follow these steps to configure the switch to receive NTP broadcast packets from connected peers:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Specify the interface to receive NTP broadcast packets, and enter interface configuration mode.
3.	ntp broadcast client	Enable the interface to receive NTP broadcast packets. By default, no interfaces receive NTP broadcast packets.
4.	exit	Return to global configuration mode.
5.	ntp broadcastdelay <i>microseconds</i>	(Optional) Change the estimated round-trip delay between the switch and the NTP broadcast server. The default is 3000 microseconds; the range is 1 to 999999.
6.	end	Return to privileged EXEC mode.
7.	show running-config	Verify your entries.
8.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable an interface from receiving NTP broadcast packets, use the **no ntp broadcast client** interface configuration command. To change the estimated round-trip delay to the default, use the **no ntp broadcastdelay** global configuration command.

This example shows how to configure an interface to receive NTP broadcast packets:

```
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# ntp broadcast client
```

Configuring NTP Access Restrictions

You can control NTP access on two levels as described in these sections:

- Creating an Access Group and Assigning a Basic IP Access List, on page 99
- Disabling NTP Services on a Specific Interface, on page 101

Creating an Access Group and Assigning a Basic IP Access List

Beginning in privileged EXEC mode, follow these steps to control access to NTP services by using access lists:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	ntp access-group { query-only serve-only serve peer } <i>access-list-number</i>	<p>Create an access group, and apply a basic IP access list.</p> <p>The keywords have these meanings:</p> <ul style="list-style-type: none"> • query-only—Allows only NTP control queries. • serve-only—Allows only time requests. • serve—Allows time requests and NTP control queries, but does not allow the switch to synchronize to the remote device. • peer—Allows time requests and NTP control queries and allows the switch to synchronize to the remote device. <p>For <i>access-list-number</i>, enter a standard IP access list number from 1 to 99.</p>
3.	access-list <i>access-list-number</i> permit <i>source</i> [<i>source-wildcard</i>]	<p>Create the access list.</p> <ul style="list-style-type: none"> • For <i>access-list-number</i>, enter the number specified in Step 2. • Enter the permit keyword to permit access if the conditions are matched. • For <i>source</i>, enter the IP address of the device that is permitted access to the switch. • (Optional) For <i>source-wildcard</i>, enter the wildcard bits to be applied to the source. <p>Note: When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.</p>
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The access group keywords are scanned in this order, from least restrictive to most restrictive:

1. **peer**—Allows time requests and NTP control queries and allows the switch to synchronize itself to a device whose address passes the access list criteria.
2. **serve**—Allows time requests and NTP control queries, but does not allow the switch to synchronize itself to a device whose address passes the access list criteria.
3. **serve-only**—Allows only time requests from a device whose address passes the access list criteria.
4. **query-only**—Allows only NTP control queries from a device whose address passes the access list criteria.

If the source IP address matches the access lists for more than one access type, the first type is granted. If no access groups are specified, all access types are granted to all devices. If any access groups are specified, only the specified access types are granted.

To remove access control to the switch NTP services, use the **no ntp access-group** {**query-only** | **serve-only** | **serve** | **peer**} global configuration command.

This example shows how to configure the switch to allow itself to synchronize to a peer from access list 99. However, the switch restricts access to allow only time requests from access list 42:

```
Switch# configure terminal
Switch(config)# ntp access-group peer 99
Switch(config)# ntp access-group serve-only 42
Switch(config)# access-list 99 permit 172.20.130.5
Switch(config)# access list 42 permit 172.20.130.6
```

Disabling NTP Services on a Specific Interface

NTP services are enabled on all interfaces by default.

Beginning in privileged EXEC mode, follow these steps to disable NTP packets from being received on an interface:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to disable.
3.	ntp disable	Disable NTP packets from being received on the interface. By default, all interfaces receive NTP packets.
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To re-enable receipt of NTP packets on an interface, use the **no ntp disable** interface configuration command.

Configuring the Source IP Address for NTP Packets

When the switch sends an NTP packet, the source IP address is normally set to the address of the interface through which the NTP packet is sent. Use the **ntp source** global configuration command when you want to use a particular source IP address for all NTP packets. The address is taken from the specified interface. This command is useful if the address on an interface cannot be used as the destination for reply packets.

Beginning in privileged EXEC mode, follow these steps to configure a specific interface from which the IP source address is to be taken:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	ntp source <i>type number</i>	Specify the interface type and number from which the IP source address is taken. By default, the source address is determined by the outgoing interface.
3.	end	Return to privileged EXEC mode.
4.	show running-config	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The specified interface is used for the source address for all packets sent to all destinations. If a source address is to be used for a specific association, use the **source** keyword in the **ntp peer** or **ntp server** global configuration command as described in the “Configuring NTP Associations” section on page 97.

Displaying the NTP Configuration

You can use two privileged EXEC commands to display NTP information:

- **show ntp associations [detail]**
- **show ntp status**

For detailed information about the fields in these displays, refer to the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference*.

Configuring Time and Date Manually

If no other source of time is available, you can manually configure the time and date after the system is restarted. The time remains accurate until the next system restart. We recommend that you use manual configuration only as a last resort. If you have an outside source to which the switch can synchronize, you do not need to manually set the system clock.

This section contains this configuration information:

- Setting the System Clock, on page 102
- Displaying the Time and Date Configuration, on page 103
- Configuring the Time Zone, on page 103
- Configuring Summer Time (Daylight Saving Time), on page 103

Setting the System Clock

If you have an outside source on the network that provides time services, such as an NTP server, you do not need to manually set the system clock.

Beginning in privileged EXEC mode, follow these steps to set the system clock:

Step	Command	Purpose
1.	clock set <i>hh:mm:ss day month year</i> or clock set <i>hh:mm:ss month day year</i>	Manually set the system clock using one of these formats. <ul style="list-style-type: none"> • For <i>hh:mm:ss</i>, specify the time in hours (24-hour format), minutes, and seconds. The time specified is relative to the configured time zone. • For <i>day</i>, specify the day by date in the month. • For <i>month</i>, specify the month by name. • For <i>year</i>, specify the year (no abbreviation).
2.	show running-config	Verify your entries.
3.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to manually set the system clock to 1:32 p.m. on July 23, 2001:

```
Switch# clock set 13:32:00 23 July 2001
```


Displaying the Time and Date Configuration

To display the time and date configuration, use the **show clock [detail]** privileged EXEC command.

The system clock keeps an *authoritative* flag that shows whether the time is authoritative (believed to be accurate). If the system clock has been set by a timing source such as NTP, the flag is set. If the time is not authoritative, it is used only for display purposes. Until the clock is authoritative and the *authoritative* flag is set, the flag prevents peers from synchronizing to the clock when the peers' time is invalid.

The symbol that precedes the **show clock** display has this meaning:

- * - Time is not authoritative.
- (blank) - Time is authoritative.
- . - Time is authoritative, but NTP is not synchronized.

Configuring the Time Zone

Beginning in privileged EXEC mode, follow these steps to manually configure the time zone:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	clock timezone <i>zone hours-offset</i> [<i>minutes-offset</i>]	Set the time zone. The switch keeps internal time in universal time coordinated (UTC), so this command is used only for display purposes and when the time is manually set. <ul style="list-style-type: none">• For <i>zone</i>, enter the name of the time zone to be displayed when standard time is in effect. The default is UTC.• For <i>hours-offset</i>, enter the hours offset from UTC.• (Optional) For <i>minutes-offset</i>, enter the minutes offset from UTC.
3.	end	Return to privileged EXEC mode.
4.	show running-config	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The *minutes-offset* variable in the **clock timezone** global configuration command is available for those cases where a local time zone is a percentage of an hour different from UTC. For example, the time zone for some sections of Atlantic Canada (AST) is UTC-3.5, where the 3 means 3 hours and .5 means 50 percent. In this case, the necessary command is **clock timezone AST -3 30**.

To set the time to UTC, use the **no clock timezone** global configuration command.

Configuring Summer Time (Daylight Saving Time)

Beginning in privileged EXEC mode, follow these steps to configure summer time (daylight saving time) in areas where it starts and ends on a particular day of the week each year:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	clock summer-time zone recurring [week day month hh:mm week day month hh:mm [offset]]	<p>Configure summer time to start and end on the specified days every year.</p> <p>Summer time is disabled by default. If you specify clock summer-time zone recurring without parameters, the summer time rules default to the United States rules.</p> <ul style="list-style-type: none"> For <i>zone</i>, specify the name of the time zone (for example, PDT) to be displayed when summer time is in effect. (Optional) For <i>week</i>, specify the week of the month (1 to 5 or last). (Optional) For <i>day</i>, specify the day of the week (Sunday, Monday...). (Optional) For <i>month</i>, specify the month (January, February...). (Optional) For <i>hh:mm</i>, specify the time (24-hour format) in hours and minutes. (Optional) For <i>offset</i>, specify the number of minutes to add during summer time. The default is 60.
3.	end	Return to privileged EXEC mode.
4.	show running-config	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The first part of the **clock summer-time** global configuration command specifies when summer time begins, and the second part specifies when it ends. All times are relative to the local time zone. The start time is relative to standard time. The end time is relative to summer time. If the starting month is after the ending month, the system assumes that you are in the southern hemisphere.

This example shows how to specify that summer time starts on the first Sunday in April at 02:00 and ends on the last Sunday in October at 02:00:

```
Switch(config)# clock summer-time PDT recurring 1 Sunday April 2:00 last Sunday
October 2:00
```

Beginning in privileged EXEC mode, follow these steps if summer time in your area does not follow a recurring pattern (configure the exact date and time of the next summer time events):

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	clock summer-time zone date [month date year hh:mm month date year hh:mm [offset]] or clock summer-time zone date [date month year hh:mm date month year hh:mm [offset]]	Configure summer time to start on the first date and end on the second date. Summer time is disabled by default. <ul style="list-style-type: none"> For <i>zone</i>, specify the name of the time zone (for example, PDT) to be displayed when summer time is in effect. (Optional) For <i>week</i>, specify the week of the month (1 to 5 or last). (Optional) For <i>day</i>, specify the day of the week (Sunday, Monday...). (Optional) For <i>month</i>, specify the month (January, February...). (Optional) For <i>hh:mm</i>, specify the time (24-hour format) in hours and minutes. (Optional) For <i>offset</i>, specify the number of minutes to add during summer time. The default is 60.
3.	end	Return to privileged EXEC mode.
4.	show running-config	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The first part of the **clock summer-time** global configuration command specifies when summer time begins, and the second part specifies when it ends. All times are relative to the local time zone. The start time is relative to standard time. The end time is relative to summer time. If the starting month is after the ending month, the system assumes that you are in the southern hemisphere.

To disable summer time, use the **no clock summer-time** global configuration command.

This example shows how to set summer time to start on October 12, 2000, at 02:00, and end on April 26, 2001, at 02:00:

```
Switch(config)# clock summer-time pdt date 12 October 2000 2:00 26 April 2001 2:00
```

Configuring a System Name and Prompt

You configure the system name on the switch to identify it. By default, the system name and prompt are *Switch*.

If you have not configured a system prompt, the first 20 characters of the system name are used as the system prompt. A greater-than symbol [**>**] is appended. The prompt is updated whenever the system name changes, unless you manually configure the prompt by using the **prompt** global configuration command.

Note: For complete syntax and usage information for the commands used in this section, refer to the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference*.

This section contains this configuration information:

- Default System Name and Prompt Configuration, on page 106
- Configuring a System Name, on page 106

- Configuring a System Prompt, on page 106
- Understanding DNS, on page 107

Default System Name and Prompt Configuration

The default switch system name and prompt is *Switch*.

Configuring a System Name

Beginning in privileged EXEC mode, follow these steps to manually configure a system name:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	hostname <i>name</i>	Manually configure a system name. The default setting is <i>switch</i> . The name must follow the rules for ARPANET host names. They must start with a letter, end with a letter or digit, and have as interior characters only letters, digits, and hyphens. Names can be up to 63 characters.
3.	end	Return to privileged EXEC mode.
4.	show running-config	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

When you set the system name, it is also used as the system prompt. You can override the prompt setting by using the **prompt** global configuration command.

To return to the default hostname, use the **no hostname** global configuration command.

Configuring a System Prompt

Beginning in privileged EXEC mode, follow these steps to manually configure a system prompt:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	prompt <i>string</i>	Configure the command-line prompt to override the setting from the hostname command. The default prompt is either <i>switch</i> or the name defined with the hostname global configuration command, followed by an angle bracket (>) for user EXEC mode or a pound sign (#) for privileged EXEC mode. The prompt can consist of all printing characters and escape sequences.
3.	end	Return to privileged EXEC mode.
4.	show running-config	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default prompt, use the **no prompt** [*string*] global configuration command.

Understanding DNS

The DNS protocol controls the Domain Name System (DNS), a distributed database with which you can map host names to IP addresses. When you configure DNS on your switch, you can substitute the host name for the IP address with all IP commands, such as **ping**, **telnet**, **connect**, and related Telnet support operations.

IP defines a hierarchical naming scheme that allows a device to be identified by its location or domain. Domain names are pieced together with periods (.) as the delimiting characters. For example, Cisco Systems is a commercial organization that IP identifies by a *com* domain name, so its domain name is *ibm.com*. A specific device in this domain, for example, the File Transfer Protocol (FTP) system is identified as *ftp.ibm.com*.

To keep track of domain names, IP has defined the concept of a domain name server, which holds a cache (or database) of names mapped to IP addresses. To map domain names to IP addresses, you must first identify the host names, specify the name server that is present on your network, and enable the DNS.

This section contains this configuration information:

- Default DNS Configuration, on page 107
- Setting Up DNS, on page 107
- Displaying the DNS Configuration, on page 108

Default DNS Configuration

Table 28 shows the default DNS configuration.

Table 28. Default DNS Configuration

Feature	Default Setting
DNS enable state	Enabled.
DNS default domain name	None configured.
DNS servers	No name server addresses are configured.

Setting Up DNS

Beginning in privileged EXEC mode, follow these steps to set up your switch to use the DNS:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	ip domain-name <i>name</i>	Define a default domain name that the software uses to complete unqualified host names (names without a dotted-decimal domain name). Do not include the initial period that separates an unqualified name from the domain name. At boot time, no domain name is configured.

Step	Command	Purpose
3.	ip name-server <i>server-address1</i> [<i>server-address2</i> ... <i>server-address6</i>]	Specify the address of one or more name servers to use for name and address resolution. You can specify up to six name servers. Separate each server address with a space. The first server specified is the primary server. The switch sends DNS queries to the primary server first. If that query fails, the backup servers are queried.
4.	ip domain-lookup	(Optional) Enable DNS-based host name-to-address translation on your switch. This feature is enabled by default. If your network devices require connectivity with devices in networks for which you do not control name assignment, you can dynamically assign device names that uniquely identify your devices by using the global Internet naming scheme (DNS).
5.	end	Return to privileged EXEC mode.
6.	show running-config	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

If you use the switch IP address as its hostname, the IP address is used and no DNS query occurs. If you configure a hostname that contains no periods (.), a period followed by the default domain name is appended to the hostname before the DNS query is made to map the name to an IP address. The default domain name is the value set by the **ip domain-name** global configuration command. If there is a period (.) in the hostname, the software looks up the IP address without appending any default domain name to the hostname.

To remove a domain name, use the **no ip domain-name** *name* global configuration command. To remove a name server address, use the **no ip name-server** *server-address* global configuration command. To disable DNS on the switch, use the **no ip domain-lookup** global configuration command.

Displaying the DNS Configuration

To display the DNS configuration information, use the **show running-config** privileged EXEC command.

Creating a Banner

You can configure a message-of-the-day (MOTD) and a login banner. The MOTD banner displays on all connected terminals at login and is useful for sending messages that affect all network users (such as impending system shutdowns).

The login banner also displays on all connected terminals. It is displayed after the MOTD banner and before the login prompts.

Note: For complete syntax and usage information for the commands used in this section, refer to the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference*.

This section contains this configuration information:

- Default Banner Configuration, on page 109
- Configuring a Message-of-the-Day Login Banner, on page 109
- Configuring a Login Banner, on page 110

Default Banner Configuration

The MOTD and login banners are not configured.

Configuring a Message-of-the-Day Login Banner

You can create a single or multiline message banner that appears on the screen when someone logs in to the switch.

Beginning in privileged EXEC mode, follow these steps to configure a MOTD login banner:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	banner motd c message c	Specify the message of the day. For <i>c</i> , enter the delimiting character of your choice, for example, a pound sign (#), and press the Return key. The delimiting character signifies the beginning and end of the banner text. Characters after the ending delimiter are discarded. For <i>message</i> , enter a banner message up to 255 characters. You cannot use the delimiting character in the message.
3.	end	Return to privileged EXEC mode.
4.	show running-config	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete the MOTD banner, use the **no banner motd** global configuration command.

This example shows how to configure a MOTD banner for the switch by using the pound sign (#) symbol as the beginning and ending delimiter:

```
Switch(config)# banner motd #  
This is a secure site. Only authorized users are allowed.  
For access, contact technical support.  
#  
Switch(config)#
```

This example shows the banner displayed from the previous configuration:

```
Unix> telnet 172.2.5.4  
Trying 172.2.5.4...  
Connected to 172.2.5.4.  
Escape character is '^]'.  
  
This is a secure site. Only authorized users are allowed.  
For access, contact technical support.
```

```
User Access Verification  
  
Password:
```

Configuring a Login Banner

You can configure a login banner to be displayed on all connected terminals. This banner appears after the MOTD banner and before the login prompt.

Beginning in privileged EXEC mode, follow these steps to configure a login banner:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	banner login c message c	Specify the login message. For <i>c</i> , enter the delimiting character of your choice, for example, a pound sign (#), and press the Return key. The delimiting character signifies the beginning and end of the banner text. Characters after the ending delimiter are discarded. For <i>message</i> , enter a login message up to 255 characters. You cannot use the delimiting character in the message.
3.	end	Return to privileged EXEC mode.
4.	show running-config	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete the login banner, use the **no banner login** global configuration command.

This example shows how to configure a login banner for the switch by using the dollar sign (\$) symbol as the beginning and ending delimiter:

```
Switch(config)# banner login $  
Access for authorized users only. Please enter your username and password.  
$  
Switch(config)#
```

Managing the MAC Address Table

The MAC address table contains address information that the switch uses to forward traffic between ports. All MAC addresses in the address table are associated with one or more ports. The address table includes these types of addresses:

- **Dynamic address**
A source MAC address that the switch learns and then ages when it is not in use.
- **Static address**
A manually entered unicast or multicast address that does not age and that is not lost when the switch resets.

The address table lists the destination MAC address, the associated VLAN ID, and port number associated with the address.

Note: For complete syntax and usage information for the commands used in this section, refer to the command reference for this release.

This section contains this configuration information:

- Building the Address Table, on page 111
- MAC Addresses and VLANs, on page 111
- Default MAC Address Table Configuration, on page 111

- Changing the Address Aging Time, on page 111
- Removing Dynamic Address Entries, on page 112
- Configuring MAC Address Notification Traps, on page 112
- Adding and Removing Static Address Entries, on page 114
- Displaying Address Table Entries, on page 115

Building the Address Table

With multiple MAC addresses supported on all ports, you can connect any port on the switch to individual workstations, repeaters, switches, routers, or other network devices. The switch provides dynamic addressing by learning the source address of packets it receives on each port and adding the address and its associated port number to the address table. As stations are added or removed from the network, the switch updates the address table, adding new dynamic addresses and aging out those that are not in use.

The aging interval is configured on a per-switch basis. However, the switch maintains an address table for each VLAN, and STP can accelerate the aging interval on a per-VLAN basis.

The switch sends packets between any combination of ports, based on the destination address of the received packet. Using the MAC address table, the switch forwards the packet only to the port or ports associated with the destination address. If the destination address is on the port that sent the packet, the packet is filtered and not forwarded. The switch always uses the store-and-forward method: complete packets are stored and checked for errors before transmission.

MAC Addresses and VLANs

All addresses are associated with a VLAN. An address can exist in more than one VLAN and have different destinations in each. Multicast addresses, for example, could be forwarded to port 1 in VLAN 1 and ports 9, 10, and 11 in VLAN 5.

Each VLAN maintains its own logical address table. A known address in one VLAN is unknown in another until it is learned or statically associated with a port in the other VLAN. Addresses that are statically entered in one VLAN must be configured as static addresses in all other VLANs or remain unlearned in the other VLANs.

Default MAC Address Table Configuration

Table 29 shows the default MAC address table configuration.

Table 29. Default MAC Address Table Configuration

Feature	Default Setting
Aging time	300 seconds
Dynamic addresses	Automatically learned
Static addresses	None configured

Changing the Address Aging Time

Dynamic addresses are source MAC addresses that the switch learns and then ages when they are not in use. The aging time parameter defines how long the switch retains unseen addresses. This parameter applies to all VLANs.

Setting too short an aging time can cause addresses to be prematurely removed from the table. Then when the switch receives a packet for an unknown destination, it floods the packet to all ports in the same VLAN as the receiving port. This unnecessary flooding can impact performance. Setting too long an aging time can cause the address table to be filled with unused addresses, which prevents new addresses from being learned.

Beginning in privileged EXEC mode, follow these steps to configure the dynamic address table aging time:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	mac address-table aging-time [0 10-1000000]	Set the length of time that a dynamic entry remains in the MAC address table after the entry is used or updated. The range is 10 to 1000000 seconds. The default is 300. You can also enter 0, which disables aging. Static address entries are never aged or removed from the table.
3.	end	Return to privileged EXEC mode.
4.	show mac address-table aging-time	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default value, use the **no mac address-table aging-time** global configuration command.

Removing Dynamic Address Entries

To remove all dynamic entries, use the **clear mac address-table dynamic** command in privileged EXEC mode. You can also remove a specific MAC address (**clear mac address-table dynamic address *mac-address***), remove all addresses on the specified physical port or port channel (**clear mac address-table dynamic interface *interface-id***), or remove all addresses on a specified VLAN (**clear mac address-table dynamic vlan *vlan-id***).

To verify that dynamic entries have been removed, use the **show mac address-table dynamic** privileged EXEC command.

Configuring MAC Address Notification Traps

MAC address notification enables you to track users on a network by storing the MAC address activity on the switch. Whenever the switch learns or removes a MAC address, an SNMP notification can be generated and sent to the NMS. If you have many users coming and going from the network, you can set a trap interval time to bundle the notification traps and reduce network traffic. The MAC notification history table stores the MAC address activity for each hardware port for which the trap is enabled. MAC address notifications are generated for dynamic and secure MAC addresses; events are not generated for self addresses, multicast addresses, or other static addresses.

Beginning in privileged EXEC mode, follow these steps to configure the switch to send MAC address notification traps to an NMS host:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	snmp-server host <i>host-addr</i> { traps / informs } { version { 1 / 2c / 3 }} <i>community-string</i> <i>notification-type</i>	Specify the recipient of the trap message. <ul style="list-style-type: none"> For <i>host-addr</i>, specify the name or address of the NMS. Specify traps (the default) to send SNMP traps to the host. Specify informs to send SNMP informs to the host. Specify the SNMP version to support. Version 1, the default, is not available with informs. For <i>community-string</i>, specify the string to send with the notification operation. Though you can set this string by using the snmp-server host command, we recommend that you define this string by using the snmp-server community command before using the snmp-server host command. For <i>notification-type</i>, use the mac-notification keyword.
3.	snmp-server enable traps mac-notification	Enable the switch to send MAC address traps to the NMS.
4.	mac address-table notification	Enable the MAC address notification feature.
5.	mac address-table notification [interval <i>value</i>] [history-size <i>value</i>]	Enter the trap interval time and the history table size. <ul style="list-style-type: none"> (Optional) For interval <i>value</i>, specify the notification trap interval in seconds between each set of traps that are generated to the NMS. The range is 0 to 2147483647 seconds; the default is 1 second. (Optional) For history-size <i>value</i>, specify the maximum number of entries in the MAC notification history table. The range is 0 to 500; the default is 1.
6.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface on which to enable the SNMP MAC address notification trap.
7.	snmp trap mac-notification { added removed }	Enable the MAC address notification trap. <ul style="list-style-type: none"> Enable the MAC notification trap whenever a MAC address is added on this interface. Enable the MAC notification trap whenever a MAC address is removed from this interface.
8.	end	Return to privileged EXEC mode.

Step	Command	Purpose
9.	show mac address-table notification interface show running-config	Verify your entries.
10.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the switch from sending MAC address notification traps, use the **no snmp-server enable traps mac-notification** global configuration command. To disable the MAC address notification traps on a specific interface, use the **no snmp trap mac-notification {added | removed}** interface configuration command. To disable the MAC address notification feature, use the **no mac address-table notification** global configuration command.

This example shows how to specify 172.20.10.10 as the NMS, enable the switch to send MAC address notification traps to the NMS, enable the MAC address notification feature, set the interval time to 60 seconds, set the history-size to 100 entries, and enable traps whenever a MAC address is added on Gigabit Ethernet interface 0/17.

```
Switch(config)# snmp-server host 172.20.10.10 traps private
Switch(config)# snmp-server enable traps mac-notification
Switch(config)# mac address-table notification
Switch(config)# mac address-table notification interval 60
Switch(config)# mac address-table notification history-size 100
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# snmp trap mac-notification added
```

You can verify the previous commands by entering the **show mac address-table notification interface** and the **show mac address-table notification** privileged EXEC commands.

Adding and Removing Static Address Entries

A static address has these characteristics:

- It is manually entered in the address table and must be manually removed.
- It can be a unicast or multicast address.
- It does not age and is retained when the switch restarts.

You can add and remove static addresses and define the forwarding behavior for them. The forwarding behavior determines how a port that receives a packet forwards it to another port for transmission. Because all ports are associated with at least one VLAN, the switch acquires the VLAN ID for the address from the ports that you specify. You can specify a different list of destination ports for each source port.

A static address in one VLAN must be a static address in other VLANs. A packet with a static address that arrives on a VLAN where it has not been statically entered is flooded to all ports and not learned.

You add a static address to the address table by specifying the destination MAC address (unicast or multicast) and the VLAN from which it is received. Packets received with this destination address are forwarded to the interface specified with the *interface-id* option.

Beginning in privileged EXEC mode, follow these steps to add a static address:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	mac address-table static <i>mac-addr</i> vlan <i>vlan-id</i> interface <i>interface-id</i>	<p>Add a static address to the MAC address table.</p> <ul style="list-style-type: none"> For <i>mac-addr</i>, specify the destination MAC address (unicast or multicast) to add to the address table. Packets with this destination address received in the specified VLAN are forwarded to the specified interface. For <i>vlan-id</i>, specify the VLAN for which the packet with the specified MAC address is received. Valid VLAN IDs are: <ul style="list-style-type: none"> VLAN ID 2 to 4094 on the internal 1000 Mbps ports VLAN ID 1 on the internal 100 Mbps management module ports VLAN ID 1 to 4094 on the external 10/100/1000 Mbps ports For <i>interface-id...</i>, specify the interface to which the received packet is forwarded. Valid interfaces include physical ports.
3.	end	Return to privileged EXEC mode.
4.	show mac address-table static	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove static entries from the address table, use the **no mac address-table static** *mac-addr* **vlan** *vlan-id* **interface** *interface-id* global configuration command.

This example shows how to add the static address c2f3.220a.12f4 to the MAC address table. When a packet is received in VLAN 4 with this MAC address as its destination address, the packets is forwarded to the specified interface:

```
Switch(config)# mac address-table static c2f3.220a.12f4 vlan 4 interface gi0/17
```

Displaying Address Table Entries

You can display the MAC address table by using one or more of the privileged EXEC commands described in Table 30:

Table 30. Commands for Displaying the MAC Address Table

Command	Description
show mac address-table address	Displays MAC address table information for the specified MAC address.
show mac address-table aging-time	Displays the aging time in all VLANs or the specified VLAN.
show mac address-table count	Displays the number of addresses present in all VLANs or the specified VLAN.
show mac address-table dynamic	Displays dynamic MAC address table entries only.
show mac address-table interface	Displays the MAC address table information for the specified interface.
show mac address-table multicast	Displays the Layer 2 multicast entries for all VLANs or the specified VLAN.

Table 30. Commands for Displaying the MAC Address Table (continued)

Command	Description
show mac address-table static	Displays static MAC address table entries only.
show mac address-table vlan	Displays the MAC address table information for the specified VLAN.

Managing the ARP Table

To communicate with a device (over Ethernet, for example), the software first must determine the 48-bit MAC or the local data link address of that device. The process of determining the local data link address from an IP address is called *address resolution*.

The Address Resolution Protocol (ARP) associates a host IP address with the corresponding media or MAC addresses and the VLAN ID. Taking an IP address as input, ARP determines the associated MAC address. Once a MAC address is determined, the IP-MAC address association is stored in an ARP cache for rapid retrieval. Then the IP datagram is encapsulated in a link-layer frame and sent over the network. Encapsulation of IP datagrams and ARP requests and replies on IEEE 802 networks other than Ethernet is specified by the Subnetwork Access Protocol (SNAP). By default, standard Ethernet-style ARP encapsulation (represented by the **arpa** keyword) is enabled on the IP interface.

ARP entries added manually to the table do not age and must be manually removed.

Chapter 7. Configuring Switch-Based Authentication

This chapter describes how to configure switch-based authentication on the switch. This chapter consists of these sections:

- Preventing Unauthorized Access to Your Switch, on page 117
- Protecting Access to Privileged EXEC Commands, on page 117
- Controlling Switch Access with TACACS+, on page 124
- Controlling Switch Access with RADIUS, on page 131
- Configuring the Switch for Local Authentication and Authorization, on page 145
- Configuring the Switch for Secure Shell, on page 146

Preventing Unauthorized Access to Your Switch

You can prevent unauthorized users from reconfiguring your switch and viewing configuration information. Typically, you want network administrators to have access to your switch while you restrict access to users who dial from outside the network through an asynchronous port, connect from outside the network through a serial port, or connect through a terminal or workstation from within the local network.

To prevent unauthorized access into your switch, you should configure one or more of these security features:

- At a minimum, you should configure passwords and privileges at each switch port. These passwords are locally stored on the switch. When users attempt to access the switch through a port or line, they must enter the password specified for the port or line before they can access the switch. For more information, see the “Protecting Access to Privileged EXEC Commands” section on page 117.
- For an additional layer of security, you can also configure username and password pairs, which are locally stored on the switch. These pairs are assigned to lines or interfaces and authenticate each user before that user can access the switch. If you have defined privilege levels, you can also assign a specific privilege level (with associated rights and privileges) to each username and password pair. For more information, see the “Configuring Username and Password Pairs” section on page 121.
- If you want to use username and password pairs, but you want to store them centrally on a server instead of locally, you can store them in a database on a security server. Multiple networking devices can then use the same database to obtain user authentication (and, if necessary, authorization) information. For more information, see the “Controlling Switch Access with TACACS+” section on page 124.

Protecting Access to Privileged EXEC Commands

A simple way of providing terminal access control in your network is to use passwords and assign privilege levels. Password protection restricts access to a network or network device. Privilege levels define what commands users can enter after they have logged into a network device.

Note: For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Security Command Reference for Cisco IOS Release 12.1*.

This section describes how to control access to the configuration file and privileged EXEC commands. It contains this configuration information:

- Default Password and Privilege Level Configuration, on page 118
- Setting or Changing a Static Enable Password, on page 118
- Protecting Enable and Enable Secret Passwords with Encryption, on page 119
- Changing a Telnet Password for a Terminal Line, on page 121
- Configuring Username and Password Pairs, on page 121
- Configuring Multiple Privilege Levels, on page 122

Default Password and Privilege Level Configuration

Table 31 shows the default password and privilege level configuration.

Table 31. Default Password and Privilege Levels

Feature	Default Setting
Enable password and privilege level	No password is defined. The default is level 15 (privileged EXEC level). The password is not encrypted in the configuration file.
Enable secret password and privilege level	No password is defined. The default is level 15 (privileged EXEC level). The password is encrypted before it is written to the configuration file.
Line password	No password is defined.

Setting or Changing a Static Enable Password

The enable password controls access to the privileged EXEC mode. Beginning in privileged EXEC mode, follow these steps to set or change a static enable password:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	enable password <i>password</i>	<p>Define a new password or change an existing password for access to privileged EXEC mode.</p> <p>By default, no password is defined.</p> <p>For <i>password</i>, specify a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. It can contain the question mark (?) character if you precede the question mark with the key combination Crtl-v when you create the password; for example, to create the password abc?123, do this:</p> <p>Enter abc.</p> <p>Enter Crtl-v.</p> <p>Enter ?123.</p> <p>When the system prompts you to enter the enable password, you need not precede the question mark with the Ctrl-v; you can simply enter abc?123 at the password prompt.</p>
3.	end	Return to privileged EXEC mode.

Step	Command	Purpose
4.	show running-config	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file. The enable password is not encrypted and can be read in the switch configuration file.

To remove the password, use the **no enable password** global configuration command.

This example shows how to change the enable password to *11u2c3k4y5*. The password is not encrypted and provides access to level 15 (traditional privileged EXEC mode access):

```
Switch(config)# enable password 11u2c3k4y5
```

Protecting Enable and Enable Secret Passwords with Encryption

To provide an additional layer of security, particularly for passwords that cross the network or that are stored on a Trivial File Transfer Protocol (TFTP) server, you can use either the **enable password** or **enable secret** global configuration commands. Both commands accomplish the same thing; that is, you can establish an encrypted password that users must enter to access privileged EXEC mode (the default) or any privilege level you specify.

We recommend that you use the **enable secret** command because it uses an improved encryption algorithm.

If you configure the **enable secret** command, it takes precedence over the **enable password** command; the two commands cannot be in effect simultaneously.

Beginning in privileged EXEC mode, follow these steps to configure encryption for enable and enable secret passwords:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	enable password [level <i>level</i>] { <i>password</i> <i>encryption-type</i> <i>encrypted-password</i> } or enable secret [level <i>level</i>] { <i>password</i> <i>encryption-type</i> <i>encrypted-password</i> }	Define a new password or change an existing password for access to privileged EXEC mode. or Define a secret password, which is saved using a nonreversible encryption method. <ul style="list-style-type: none"> (Optional) For <i>level</i>, the range is from 0 to 15. Level 1 is normal user EXEC mode privileges. The default level is 15 (privileged EXEC mode privileges). For <i>password</i>, specify a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. By default, no password is defined. (Optional) For <i>encryption-type</i>, only type 5, a Cisco proprietary encryption algorithm, is available. If you specify an encryption type, you must provide an encrypted password—an encrypted password you copy from switch configuration. <p>Note: If you specify an encryption type and then enter a clear text password, you can not re-enter privileged EXEC mode. You cannot recover a lost encrypted password by any method.</p>
3.	service password-encryption	(Optional) Encrypt the password when the password is defined or when the configuration is written. Encryption prevents the password from being readable in the configuration file.
4.	end	Return to privileged EXEC mode.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

If both the enable and enable secret passwords are defined, users must enter the enable secret password.

Use the **level** keyword to define a password for a specific privilege level. After you specify the level and set a password, give the password only to users who need to have access at this level. Use the **privilege level** global configuration command to specify commands accessible at various levels. For more information, see the “Configuring Multiple Privilege Levels” section on page 122.

If you enable password encryption, it applies to all passwords including username passwords, authentication key passwords, the privileged command password, and virtual terminal line password.

To remove a password and level, use the **no enable password** [**level** *level*] or **no enable secret** [**level** *level*] global configuration command. To disable password encryption, use the **no service password-encryption** global configuration command.

This example shows how to configure the encrypted password `1FaD0$Xyti5Rkls3LoyxzS8` for privilege level 2:

```
Switch(config)# enable secret level 2 5 $1$FaD0$Xyti5Rk1s3LoyxzS8
```

Changing a Telnet Password for a Terminal Line

Note: The switch has a default username and password, which are required when accessing the switch through a Telnet session. For more information, refer to the *Cisco Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter Installation Guide*.

Beginning in privileged EXEC mode, follow these steps to configure your switch for Telnet access:

Step	Command	Purpose
1.	enable password <i>password</i>	Enter privileged EXEC mode. Note: An enable password is configured by default. It might not be necessary to a password to enter privileged EXEC mode.
2.	configure terminal	Enter global configuration mode.
3.	line vty 0 15	Configure the number of Telnet sessions (lines), and enter line configuration mode. The default configuration is login local . There are 16 possible sessions on a command-capable switch. The 0 and 15 mean that you are configuring all 16 possible Telnet sessions.
4.	password <i>password</i>	Enter a Telnet password for the line or lines. For <i>password</i> , specify a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. By default, no password is defined.
5.	end	Return to privileged EXEC mode.
6.	show running-config	Verify your entries. The password is listed under the command line vty 0 15 .
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the password, use the **no password** global configuration command.

This example shows how to set the Telnet password to *let45me67in89*:

```
Switch(config)# line vty 10  
Switch(config-line)# password let45me67in89
```

Configuring Username and Password Pairs

You can configure username and password pairs, which are locally stored on the switch. These pairs are assigned to lines or interfaces and authenticate each user before that user can access the switch. If you have defined privilege levels, you can also assign a specific privilege level (with associated rights and privileges) to each username and password pair.

Beginning in privileged EXEC mode, follow these steps to establish a username-based authentication system that requests a login username and a password:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	username <i>name</i> [privilege <i>level</i>] { password <i>encryption-type</i> <i>password</i> }	Enter the username, privilege level, and password for each user. <ul style="list-style-type: none"> For <i>name</i>, specify the user ID as one word. Spaces and quotation marks are not allowed. (Optional) For <i>level</i>, specify the privilege level the user has after gaining access. The range is 0 to 15. Level 15 gives privileged EXEC mode access. Level 1 gives user EXEC mode access. For <i>encryption-type</i>, enter 0 to specify that an unencrypted password will follow. Enter 7 to specify that a hidden password will follow. For <i>password</i>, specify the password the user must enter to gain access to the switch. The password must be from 1 to 25 characters, can contain embedded spaces, and must be the last option specified in the username command.
3.	line console 0 or line vty 0 15	Enter line configuration mode, and configure the VTY line (line 0 to 15).
4.	login local	Enable local password checking at login time. Authentication is based on the username specified in Step 2.
5.	end	Return to privileged EXEC mode.
6.	show running-config	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable username authentication for a specific user, use the **no username** *name* global configuration command. To disable password checking and allow connections without a password, use the **no login** line configuration command.

Configuring Multiple Privilege Levels

By default, the software has two modes of password security: user EXEC and privileged EXEC. You can configure up to 16 hierarchical levels of commands for each mode. By configuring multiple passwords, you can allow different sets of users to have access to specified commands.

For example, if you want many users to have access to the **clear line** command, you can assign it level 2 security and distribute the level 2 password fairly widely. But if you want more restricted access to the **configure** command, you can assign it level 3 security and distribute that password to a more restricted group of users.

This section includes this configuration information:

- Setting the Privilege Level for a Command, on page 123
- Changing the Default Privilege Level for Lines, on page 123
- Logging into and Exiting a Privilege Level, on page 124

Setting the Privilege Level for a Command

Beginning in privileged EXEC mode, follow these steps to set the privilege level for a command mode:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	privilege mode level level command	Set the privilege level for a command. <ul style="list-style-type: none"> For <i>mode</i>, enter configure for global configuration mode, exec for EXEC mode, interface for interface configuration mode, or line for line configuration mode. For <i>level</i>, the range is from 0 to 15. Level 1 is for normal user EXEC mode privileges. Level 15 is the level of access permitted by the enable password. For <i>command</i>, specify the command to which you want to restrict access.
3.	enable password level level password	Specify the enable password for the privilege level. <ul style="list-style-type: none"> For <i>level</i>, the range is from 0 to 15. Level 1 is for normal user EXEC mode privileges. For <i>password</i>, specify a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. By default, no password is defined.
4.	end	Return to privileged EXEC mode.
5.	show running-config or show privilege	Verify your entries. The first command displays the password and access level configuration. The second command displays the privilege level configuration.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

When you set a command to a privilege level, all commands whose syntax is a subset of that command are also set to that level. For example, if you set the **show ip traffic** command to level 15, the **show** commands and **show ip** commands are automatically set to privilege level 15 unless you set them individually to different levels.

To return to the default privilege for a given command, use the **no privilege mode level level command** global configuration command.

This example shows how to set the **configure** command to privilege level 14 and define *SecretPswd14* as the password users must enter to use level 14 commands:

```
Switch(config)# privilege exec level 14 configure
Switch(config)# enable password level 14 SecretPswd14
```

Changing the Default Privilege Level for Lines

Beginning in privileged EXEC mode, follow these steps to change the default privilege level for a line:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	line vty line	Select the virtual terminal line on which to restrict access.

Step	Command	Purpose
3.	privilege level <i>level</i>	Change the default privilege level for the line. For <i>level</i> , the range is from 0 to 15. Level 1 is for normal user EXEC mode privileges. Level 15 is the level of access permitted by the enable password.
4.	end	Return to privileged EXEC mode.
5.	show running-config or show privilege	Verify your entries. The first command displays the password and access level configuration. The second command displays the privilege level configuration.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Users can override the privilege level you set using the **privilege level** line configuration command by logging in to the line and enabling a different privilege level. They can lower the privilege level by using the **disable** command. If users know the password to a higher privilege level, they can use that password to enable the higher privilege level.

To return to the default line privilege level, use the **no privilege level** line configuration command.

Logging into and Exiting a Privilege Level

Beginning in privileged EXEC mode, follow these steps to log in to a specified privilege level and to exit to a specified privilege level:

Step	Command	Purpose
1.	enable <i>level</i>	Log in to a specified privilege level. For <i>level</i> , the range is 0 to 15.
2.	disable <i>level</i>	Exit to a specified privilege level. For <i>level</i> , the range is 0 to 15.

Controlling Switch Access with TACACS+

This section describes how to enable and configure Terminal Access Controller Access Control System Plus (TACACS+), which provides detailed accounting information and flexible administrative control over authentication and authorization processes. TACACS+ is facilitated through authentication, authorization, accounting (AAA) and can be enabled only through AAA commands.

Note: For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Security Command Reference for Cisco IOS Release 12.1*.

This section contains this configuration information:

- Understanding TACACS+, on page 125
- TACACS+ Operation, on page 126
- Configuring TACACS+, on page 127
- Displaying the TACACS+ Configuration, on page 131


Understanding TACACS+

TACACS+ is a security application that provides centralized validation of users attempting to gain access to your switch. TACACS+ services are maintained in a database on a TACACS+ daemon typically running on a UNIX or Windows NT workstation. You should have access to and should configure a TACACS+ server before the configuring TACACS+ features on your switch.

TACACS+ provides for separate and modular authentication, authorization, and accounting facilities. TACACS+ allows for a single access control server (the TACACS+ daemon) to provide each service—authentication, authorization, and accounting—independently. Each service can be tied into its own database to take advantage of other services available on that server or on the network, depending on the capabilities of the daemon.

The goal of TACACS+ is to provide a method for managing multiple network access points from a single management service. Your switch can be a network access server along with other Cisco routers and access servers. A network access server provides connections to a single user, to a network or subnetwork, and to interconnected networks as shown in Figure 29

Figure 29. Typical TACACS+ Network Configuration



TACACS+, administered through the AAA security services, can provide these services:

- Authentication—Provides complete control of authentication through login and password dialog, challenge and response, and messaging support.

The authentication facility can conduct a dialog with the user (for example, after a username and password are provided, to challenge a user with several questions, such as home address, mother's maiden name, service type, and social security

number). The TACACS+ authentication service can also send messages to user screens. For example, a message could notify users that their passwords must be changed because of the company's password aging policy.

- Authorization—Provides fine-grained control over user capabilities for the duration of the user's session, including but not limited to setting autocommands, access control, session duration, or protocol support. You can also enforce restrictions on what commands a user can execute with the TACACS+ authorization feature.
- Accounting—Collects and sends information used for billing, auditing, and reporting to the TACACS+ daemon. Network managers can use the accounting facility to track user activity for a security audit or to provide information for user billing. Accounting records include user identities, start and stop times, executed commands (such as PPP), number of packets, and number of bytes.

The TACACS+ protocol provides authentication between the switch and the TACACS+ daemon, and it ensures confidentiality because all protocol exchanges between the switch and the TACACS+ daemon are encrypted.

You need a system running the TACACS+ daemon software to use TACACS+ on your switch.

TACACS+ Operation

When a user attempts a simple ASCII login by authenticating to a switch using TACACS+, this process occurs:

1. When the connection is established, the switch contacts the TACACS+ daemon to obtain a username prompt, which is then displayed to the user. The user enters a username, and the switch then contacts the TACACS+ daemon to obtain a password prompt. The switch displays the password prompt to the user, the user enters a password, and the password is then sent to the TACACS+ daemon.

TACACS+ allows a conversation to be held between the daemon and the user until the daemon receives enough information to authenticate the user. The daemon prompts for a username and password combination, but can include other items, such as the user's mother's maiden name.

2. The switch eventually receives one of these responses from the TACACS+ daemon:
 - ACCEPT—The user is authenticated and service can begin. If the switch is configured to require authorization, authorization begins at this time.
 - REJECT—The user is not authenticated. The user can be denied access or is prompted to retry the login sequence, depending on the TACACS+ daemon.
 - ERROR—An error occurred at some time during authentication with the daemon or in the network connection between the daemon and the switch. If an ERROR response is received, the switch typically tries to use an alternative method for authenticating the user.
 - CONTINUE—The user is prompted for additional authentication information.

After authentication, the user undergoes an additional authorization phase if authorization has been enabled on the switch. Users must first successfully complete TACACS+ authentication before proceeding to TACACS+ authorization.

3. If TACACS+ authorization is required, the TACACS+ daemon is again contacted, and it returns an ACCEPT or REJECT authorization response. If an ACCEPT response is returned, the response contains data in the form of attributes that direct the EXEC or NETWORK session for that user, determining the services that the user can access:
 - Telnet, Secure Shell (SSH), rlogin, or privileged EXEC services

- Connection parameters, including the host or client IP address, access list, and user timeouts

Configuring TACACS+

This section describes how to configure your switch to support TACACS+. At a minimum, you must identify the host or hosts maintaining the TACACS+ daemon and define the method lists for TACACS+ authentication. You can optionally define method lists for TACACS+ authorization and accounting. A method list defines the sequence and methods to be used to authenticate, to authorize, or to keep accounts on a user. You can use method lists to designate one or more security protocols to be used, thus ensuring a backup system if the initial method fails. The software uses the first method listed to authenticate, to authorize, or to keep accounts on users; if that method does not respond, the software selects the next method in the list. This process continues until there is successful communication with a listed method or the method list is exhausted.

This section contains this configuration information:

- Default TACACS+ Configuration, on page 127
- Identifying the TACACS+ Server Host and Setting the Authentication Key, on page 127
- Configuring TACACS+ Login Authentication, on page 128
- Configuring TACACS+ Authorization for Privileged EXEC Access and Network Services, on page 130
- Starting TACACS+ Accounting, on page 131

Default TACACS+ Configuration

TACACS+ and AAA are disabled by default.

To prevent a lapse in security, you cannot configure TACACS+ through a network management application. When enabled, TACACS+ can authenticate users accessing the switch through the CLI.

Note: Although TACACS+ configuration is performed through the CLI, the TACACS+ server authenticates HTTP connections that have been configured with a privilege level of 15.

Identifying the TACACS+ Server Host and Setting the Authentication Key

You can configure the switch to use a single server or AAA server groups to group existing server hosts for authentication. You can group servers to select a subset of the configured server hosts and use them for a particular service. The server group is used with a global server-host list and contains the list of IP addresses of the selected server hosts.

Beginning in privileged EXEC mode, follow these steps to identify the IP host or host maintaining TACACS+ server and optionally set the encryption key:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	tacacs-server host <i>hostname</i> [port <i>integer</i>] [timeout <i>integer</i>] [key <i>string</i>]	Identify the IP host or hosts maintaining a TACACS+ server. Enter this command multiple times to create a list of preferred hosts. The software searches for hosts in the order in which you specify them. <ul style="list-style-type: none"> For <i>hostname</i>, specify the name or IP address of the host. (Optional) For port <i>integer</i>, specify a server port number. The default is port 49. The range is 1 to 65535. (Optional) For timeout <i>integer</i>, specify a time in seconds the switch waits for a response from the daemon before it times out and declares an error. The default is 5 seconds. The range is 1 to 1000 seconds. (Optional) For key <i>string</i>, specify the encryption key for encrypting and decrypting all traffic between the switch and the TACACS+ daemon. You must configure the same key on the TACACS+ daemon for encryption to be successful.
3.	aaa new-model	Enable AAA.
4.	aaa group server tacacs+ <i>group-name</i>	(Optional) Define the AAA server-group with a group name. This command puts the switch in a server group subconfiguration mode.
5.	server <i>ip-address</i>	(Optional) Associate a particular TACACS+ server with the defined server group. Repeat this step for each TACACS+ server in the AAA server group. Each server in the group must be previously defined in Step 2.
6.	end	Return to privileged EXEC mode.
7.	show tacacs	Verify your entries.
8.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the specified TACACS+ server name or address, use the **no tacacs-server host** *hostname* global configuration command. To remove a server group from the configuration list, use the **no aaa group server tacacs+** *group-name* global configuration command. To remove the IP address of a TACACS+ server, use the **no server ip-address** server group subconfiguration command.

Configuring TACACS+ Login Authentication

To configure AAA authentication, you define a named list of authentication methods and then apply that list to various interfaces. The method list defines the types of authentication to be performed and the sequence in which they are performed; it must be applied to a specific interface before any of the defined authentication methods are performed. The only exception is the default method list (which, by coincidence, is named *default*). The default method list is automatically applied to all interfaces except those that have a named method list explicitly defined. A defined method list overrides the default method list.

A method list describes the sequence and authentication methods to be queried to authenticate a user. You can designate one or more security protocols to be used for authentication, thus ensuring a backup system for authentication in case the initial method fails. The software uses the first method listed to authenticate users; if that method fails to respond, the software selects the next authentication method in the

method list. This process continues until there is successful communication with a listed authentication method or until all defined methods are exhausted. If authentication fails at any point in this cycle—meaning that the security server or local username database responds by denying the user access—the authentication process stops, and no other authentication methods are attempted.

Beginning in privileged EXEC mode, follow these steps to configure login authentication:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	aaa new-model	Enable AAA.
3.	aaa authentication login { default <i>list-name</i> } <i>method1</i> [<i>method2...</i>]	<p>Create a login authentication method list.</p> <ul style="list-style-type: none"> To create a default list that is used when a named list is <i>not</i> specified in the login authentication command, use the default keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all interfaces. For <i>list-name</i>, specify a character string to name the list you are creating. For <i>method1...</i>, specify the actual method the authentication algorithm tries. The additional methods of authentication are used only if the previous method returns an error, not if it fails. <p>Select one of these methods:</p> <ul style="list-style-type: none"> enable—Use the enable password for authentication. Before you can use this authentication method, you must define an enable password by using the enable password global configuration command. group tacacs+—Uses TACACS+ authentication. Before you can use this authentication method, you must configure the TACACS+ server. For more information, see the “Identifying the TACACS+ Server Host and Setting the Authentication Key” section on page 127. line—Use the line password for authentication. Before you can use this authentication method, you must define a line password. Use the password password line configuration command. local—Use the local username database for authentication. You must enter username information in the database. Use the username password global configuration command. local-case—Use a case-sensitive local username database for authentication. You must enter username information in the database by using the username name password global configuration command. none—Do not use any authentication for login.
4.	line [console tty vty] <i>line-number</i> [<i>ending-line-number</i>]	Enter line configuration mode, and configure the lines to which you want to apply the authentication list.

Step	Command	Purpose
5.	login authentication {default <i>list-name</i> }	Apply the authentication list to a line or set of lines. <ul style="list-style-type: none"> If you specify default, use the default list created with the aaa authentication login command. For <i>list-name</i>, specify the list created with the aaa authentication login command.
6.	end	Return to privileged EXEC mode.
7.	show running-config	Verify your entries.
8.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable AAA, use the **no aaa new-model** global configuration command. To disable AAA authentication, use the **no aaa authentication login** {default | *list-name*} *method1* [*method2*...] global configuration command. To either disable TACACS+ authentication for logins or to return to the default value, use the **no login authentication** {default | *list-name*} line configuration command.

Configuring TACACS+ Authorization for Privileged EXEC Access and Network Services

AAA authorization limits the services available to a user. When AAA authorization is enabled, the switch uses information retrieved from the user's profile, which is located either in the local user database or on the security server, to configure the user's session. The user is granted access to a requested service only if the information in the user profile allows it.

You can use the **aaa authorization** global configuration command with the **tacacs+** keyword to set parameters that restrict a user's network access to privileged EXEC mode.

The **aaa authorization exec tacacs+ local** command sets these authorization parameters:

- Use TACACS+ for privileged EXEC access authorization if authentication was performed by using TACACS+.
- Use the local database if authentication was not performed by using TACACS+.

Note: Authorization is bypassed for authenticated users who log in through the CLI even if authorization has been configured.

Beginning in privileged EXEC mode, follow these steps to specify TACACS+ authorization for privileged EXEC access and network services:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	aaa authorization network tacacs+	Configure the switch for user TACACS+ authorization for all network-related service requests.
3.	aaa authorization exec tacacs+	Configure the switch for user TACACS+ authorization to determine if the user has privileged EXEC access. The exec keyword might return user profile information (such as autocommand information).
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable authorization, use the **no aaa authorization {network | exec} method1** global configuration command.

Starting TACACS+ Accounting

The AAA accounting feature tracks the services that users are accessing and the amount of network resources that they are consuming. When AAA accounting is enabled, the switch reports user activity to the TACACS+ security server in the form of accounting records. Each accounting record contains accounting attribute-value (AV) pairs and is stored on the security server. This data can then be analyzed for network management, client billing, or auditing.

Beginning in privileged EXEC mode, follow these steps to enable TACACS+ accounting for each Cisco IOS privilege level and for network services:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	aaa accounting network start-stop tacacs+	Enable TACACS+ accounting for all network-related service requests.
3.	aaa accounting exec start-stop tacacs+	Enable TACACS+ accounting to send a start-record accounting notice at the beginning of a privileged EXEC process and a stop-record at the end.
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable accounting, use the **no aaa accounting {network | exec} {start-stop} method1...** global configuration command.

Displaying the TACACS+ Configuration

To display TACACS+ server statistics, use the **show tacacs** privileged EXEC command.

Controlling Switch Access with RADIUS

This section describes how to enable and configure the Remote Authentication Dial-In User Service (RADIUS), which provides detailed accounting information and flexible administrative control over authentication and authorization processes. RADIUS is facilitated through AAA and can be enabled only through AAA commands.

Note: For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Security Command Reference for Cisco IOS Release 12.1*.

This section contains this configuration information:

- Understanding RADIUS, on page 132
- RADIUS Operation, on page 133
- Configuring RADIUS, on page 133
- Displaying the RADIUS Configuration, on page 145

Understanding RADIUS

RADIUS is a distributed client/server system that secures networks against unauthorized access. RADIUS clients run on supported Cisco routers and switches. Clients send authentication requests to a central RADIUS server, which contains all user authentication and network service access information. The RADIUS host is normally a multiuser system running RADIUS server software from Cisco (Cisco Secure Access Control Server version 3.0), Livingston, Merit, Microsoft, or another software provider. For more information, refer to the RADIUS server documentation.


Use RADIUS in these network environments that require access security:

- Networks with multiple-vendor access servers, each supporting RADIUS. For example, access servers from several vendors use a single RADIUS server-based security database. In an IP-based network with multiple vendors' access servers, dial-in users are authenticated through a RADIUS server that has been customized to work with the Kerberos security system.
- Turnkey network security environments in which applications support the RADIUS protocol, such as in an access environment that uses a *smart card* access control system. In one case, RADIUS has been used with Enigma's security cards to validate users and to grant access to network resources.
- Networks already using RADIUS. You can add a Cisco switch containing a RADIUS client to the network. This might be the first step when you make a transition to a TACACS+ server. See Figure 30. on page 133.
- Network in which the user must only access a single service. Using RADIUS, you can control user access to a single host, to a single utility such as Telnet, or to the network through a protocol such as IEEE 802.1X. For more information about this protocol, see Chapter 8 "Configuring 802.1X Port-Based Authentication."
- Networks that require resource accounting. You can use RADIUS accounting independently of RADIUS authentication or authorization. The RADIUS accounting functions allow data to be sent at the start and end of services, showing the amount of resources (such as time, packets, bytes, and so forth) used during the session. An Internet service provider might use a freeware-based version of RADIUS access control and accounting software to meet special security and billing needs.

RADIUS is not suitable in these network security situations:

- Multiprotocol access environments. RADIUS does not support AppleTalk Remote Access (ARA), NetBIOS Frame Control Protocol (NBFCP), NetWare Asynchronous Services Interface (NASI), or X.25 PAD connections.
- Switch-to-switch or router-to-router situations. RADIUS does not provide two-way authentication. RADIUS can be used to authenticate from one device to a non-Cisco device if the non-Cisco device requires authentication.
- Networks using a variety of services. RADIUS generally binds a user to one service model.

Figure 30. Transitioning from RADIUS to TACACS+ Services



74721

RADIUS Operation

When a user attempts to log in and authenticate to a switch that is access controlled by a RADIUS server, these events occur:

1. The user is prompted to enter a username and password.
2. The username and encrypted password are sent over the network to the RADIUS server.
3. The user receives one of these responses from the RADIUS server:
 - ACCEPT—The user is authenticated.
 - REJECT—The user is either not authenticated and is prompted to re-enter the username and password, or access is denied.
 - CHALLENGE—A challenge requires additional data from the user.
 - CHALLENGE PASSWORD—A response requests the user to select a new password.

The ACCEPT or REJECT response is bundled with additional data that is used for privileged EXEC or network authorization. Users must first successfully complete RADIUS authentication before proceeding to RADIUS authorization, if it is enabled. The additional data included with the ACCEPT or REJECT packets includes these items:

- Telnet, SSH, rlogin, or privileged EXEC services
- Connection parameters, including the host or client IP address, access list, and user timeouts

Configuring RADIUS

This section describes how to configure your switch to support RADIUS. At a minimum, you must identify the host or hosts that run the RADIUS server software and define the method lists for RADIUS authentication. You can optionally define method lists for RADIUS authorization and accounting.

A method list defines the sequence and methods to be used to authenticate, to authorize, or to keep accounts on a user. You can use method lists to designate one or more security protocols to be used (such as TACACS+ or local username lookup),

thus ensuring a backup system if the initial method fails. The software uses the first method listed to authenticate, to authorize, or to keep accounts on users; if that method does not respond, the software selects the next method in the list. This process continues until there is successful communication with a listed method or the method list is exhausted.

You should have access to and should configure a RADIUS server before configuring RADIUS features on your switch.

This section contains this configuration information:

- Default RADIUS Configuration, on page 134
- Identifying the RADIUS Server Host, on page 134 (required)
- Configuring RADIUS Login Authentication, on page 137 (required)
- Defining AAA Server Groups, on page 139 (optional)
- Configuring RADIUS Authorization for User Privileged Access and Network Services, on page 141 (optional)
- Starting RADIUS Accounting, on page 142 (optional)
- Configuring Settings for All RADIUS Servers, on page 142 (optional)
- Configuring the Switch to Use Vendor-Specific RADIUS Attributes, on page 143 (optional)
- Configuring the Switch for Vendor-Proprietary RADIUS Server Communication, on page 144 (optional)

Default RADIUS Configuration

RADIUS and AAA are disabled by default.

To prevent a lapse in security, you cannot configure RADIUS through a network management application. When enabled, RADIUS can authenticate users accessing the switch through the CLI.

Identifying the RADIUS Server Host

Switch-to-RADIUS-server communication involves several components:

- Host name or IP address
- Authentication destination port
- Accounting destination port
- Key string
- Timeout period
- Retransmission value

You identify RADIUS security servers by their host name or IP address, host name and specific UDP port numbers, or their IP address and specific UDP port numbers. The combination of the IP address and the UDP port number creates a unique identifier, allowing different ports to be individually defined as RADIUS hosts providing a specific AAA service. This unique identifier enables RADIUS requests to be sent to multiple UDP ports on a server at the same IP address.

If two different host entries on the same RADIUS server are configured for the same service—for example, accounting—the second host entry configured acts as a fail-over backup to the first one. Using this example, if the first host entry fails to provide accounting services, the switch tries the second host entry configured on the same

device for accounting services. (The RADIUS host entries are tried in the order that they are configured.)

A RADIUS server and the switch use a shared secret text string to encrypt passwords and exchange responses. To configure RADIUS to use the AAA security commands, you must specify the host running the RADIUS server daemon and a secret text (key) string that it shares with the switch.

The timeout, retransmission, and encryption key values can be configured globally for all RADIUS servers, on a per-server basis, or in some combination of global and per-server settings. To apply these settings globally to all RADIUS servers communicating with the switch, use the three unique global configuration commands: **radius-server timeout**, **radius-server retransmit**, and **radius-server key**. To apply these values on a specific RADIUS server, use the **radius-server host** global configuration command.

Note: If you configure both global and per-server functions (timeout, retransmission, and key commands) on the switch, the per-server timer, retransmission, and key value commands override global timer, retransmission, and key value commands. For information on configuring these setting on all RADIUS servers, see the “Configuring Settings for All RADIUS Servers” section on page 142.

You can configure the switch to use AAA server groups to group existing server hosts for authentication. For more information, see the “Defining AAA Server Groups” section on page 139.

Beginning in privileged EXEC mode, follow these steps to configure per-server RADIUS server communication. This procedure is required.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	radius-server host { <i>hostname</i> <i>ip-address</i> } [auth-port <i>port-number</i>] [acct-port <i>port-number</i>] [timeout <i>seconds</i>] [retransmit <i>retries</i>] [key <i>string</i>]	<p>Specify the IP address or host name of the remote RADIUS server host.</p> <ul style="list-style-type: none"> (Optional) For auth-port <i>port-number</i>, specify the UDP destination port for authentication requests. (Optional) For acct-port <i>port-number</i>, specify the UDP destination port for accounting requests. (Optional) For timeout <i>seconds</i>, specify the time interval that the switch waits for the RADIUS server to reply before resending. The range is 1 to 1000. This setting overrides the radius-server timeout global configuration command setting. If no timeout is set with the radius-server host command, the setting of the radius-server timeout command is used. (Optional) For retransmit <i>retries</i>, specify the number of times a RADIUS request is resent to a server if that server is not responding or responding slowly. The range is 1 to 1000. If no retransmit value is set with the radius-server host command, the setting of the radius-server retransmit global configuration command is used. (Optional) For key <i>string</i>, specify the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. <p>Note: The key is a text string that must match the encryption key used on the RADIUS server. Always configure the key as the last item in the radius-server host command. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.</p> <p>To configure the switch to recognize more than one host entry associated with a single IP address, enter this command as many times as necessary, making sure that each UDP port number is different. The switch software searches for hosts in the order in which you specify them. Set the timeout, retransmit, and encryption key values to use with the specific RADIUS host.</p>
3.	end	Return to privileged EXEC mode.
4.	show running-config	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the specified RADIUS server, use the **no radius-server host** *hostname* | *ip-address* global configuration command.

This example shows how to configure one RADIUS server to be used for authentication and another to be used for accounting:

```
Switch(config)# radius-server host 172.29.36.49 auth-port 1612 key rad1
```

```
Switch(config)# radius-server host 172.20.36.50 acct-port 1618 key rad2
```

This example shows how to configure *host1* as the RADIUS server and to use the default ports for both authentication and accounting:

```
Switch(config)# radius-server host host1
```

Note: You also need to configure some settings on the RADIUS server. These settings include the IP address of the switch and the key string to be shared by both the server and the switch. For more information, refer to the RADIUS server documentation.

Configuring RADIUS Login Authentication

To configure AAA authentication, you define a named list of authentication methods and then apply that list to various interfaces. The method list defines the types of authentication to be performed and the sequence in which they are performed; it must be applied to a specific interface before any of the defined authentication methods are performed. The only exception is the default method list (which, by coincidence, is named *default*). The default method list is automatically applied to all interfaces except those that have a named method list explicitly defined.

A method list describes the sequence and authentication methods to be queried to authenticate a user. You can designate one or more security protocols to be used for authentication, thus ensuring a backup system for authentication in case the initial method fails. The software uses the first method listed to authenticate users; if that method fails to respond, the software selects the next authentication method in the method list. This process continues until there is successful communication with a listed authentication method or until all defined methods are exhausted. If authentication fails at any point in this cycle—meaning that the security server or local username database responds by denying the user access—the authentication process stops, and no other authentication methods are attempted.

Beginning in privileged EXEC mode, follow these steps to configure login authentication. This procedure is required.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	aaa new-model	Enable AAA.

Step	Command	Purpose
3.	aaa authentication login { default <i>list-name</i> } <i>method1</i> [<i>method2...</i>]	<p>Create a login authentication method list.</p> <ul style="list-style-type: none"> To create a default list that is used when a named list is <i>not</i> specified in the login authentication command, use the default keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all interfaces. For <i>list-name</i>, specify a character string to name the list you are creating. For <i>method1...</i>, specify the actual method the authentication algorithm tries. The additional methods of authentication are used only if the previous method returns an error, not if it fails. <p>Select one of these methods:</p> <ul style="list-style-type: none"> enable—Use the enable password for authentication. Before you can use this authentication method, you must define an enable password by using the enable password global configuration command. group radius—Use RADIUS authentication. Before you can use this authentication method, you must configure the RADIUS server. For more information, see the “Identifying the RADIUS Server Host” section on page 134. line—Use the line password for authentication. Before you can use this authentication method, you must define a line password. Use the password password line configuration command. local—Use the local username database for authentication. You must enter username information in the database. Use the username name password global configuration command. local-case—Use a case-sensitive local username database for authentication. You must enter username information in the database by using the username password global configuration command. none—Do not use any authentication for login.
4.	line [console tty vty] <i>line-number</i> [<i>ending-line-number</i>]	Enter line configuration mode, and configure the lines to which you want to apply the authentication list.
5.	login authentication { default <i>list-name</i> }	<p>Apply the authentication list to a line or set of lines.</p> <ul style="list-style-type: none"> If you specify default, use the default list created with the aaa authentication login command. For <i>list-name</i>, specify the list created with the aaa authentication login command.
6.	end	Return to privileged EXEC mode.
7.	show running-config	Verify your entries.
8.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable AAA, use the **no aaa new-model** global configuration command. To disable AAA authentication, use the **no aaa authentication login** {**default** | *list-name*} *method1* [*method2...*] global configuration command. To either disable RADIUS

authentication for logins or to return to the default value, use the **no login authentication {default | list-name}** line configuration command.

Defining AAA Server Groups

You can configure the switch to use AAA server groups to group existing server hosts for authentication. You select a subset of the configured server hosts and use them for a particular service. The server group is used with a global server-host list, which lists the IP addresses of the selected server hosts.

Server groups also can include multiple host entries for the same server if each entry has a unique identifier (the combination of the IP address and UDP port number), allowing different ports to be individually defined as RADIUS hosts providing a specific AAA service. If you configure two different host entries on the same RADIUS server for the same service, (for example, accounting), the second configured host entry acts as a fail-over backup to the first one.

You use the **server** group server configuration command to associate a particular server with a defined group server. You can either identify the server by its IP address or identify multiple host instances or entries by using the optional **auth-port** and **acct-port** keywords.

Beginning in privileged EXEC mode, follow these steps to define the AAA server group and associate a particular RADIUS server with it:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	radius-server host { <i>hostname</i> <i>ip-address</i> } [auth-port <i>port-number</i>] [acct-port <i>port-number</i>] [timeout <i>seconds</i>] [retransmit <i>retries</i>] [key <i>string</i>]	<p>Specify the IP address or host name of the remote RADIUS server host.</p> <ul style="list-style-type: none"> (Optional) For auth-port <i>port-number</i>, specify the UDP destination port for authentication requests. (Optional) For acct-port <i>port-number</i>, specify the UDP destination port for accounting requests. (Optional) For timeout <i>seconds</i>, specify the time interval that the switch waits for the RADIUS server to reply before resending. The range is 1 to 1000. This setting overrides the radius-server timeout global configuration command setting. If no timeout is set with the radius-server host command, the setting of the radius-server timeout command is used. (Optional) For retransmit <i>retries</i>, specify the number of times a RADIUS request is resent to a server if that server is not responding or responding slowly. The range is 1 to 1000. If no retransmit value is set with the radius-server host command, the setting of the radius-server retransmit global configuration command is used. (Optional) For key <i>string</i>, specify the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. <p>Note: The key is a text string that must match the encryption key used on the RADIUS server. Always configure the key as the last item in the radius-server host command. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.</p> <p>To configure the switch to recognize more than one host entry associated with a single IP address, enter this command as many times as necessary, making sure that each UDP port number is different. The switch software searches for hosts in the order in which you specify them. Set the timeout, retransmit, and encryption key values to use with the specific RADIUS host.</p>
3.	aaa new-model	Enable AAA.
4.	aaa group server radius <i>group-name</i>	<p>Define the AAA server-group with a group name.</p> <p>This command puts the switch in a server group configuration mode.</p>
5.	server <i>ip-address</i>	<p>Associate a particular RADIUS server with the defined server group. Repeat this step for each RADIUS server in the AAA server group.</p> <p>Each server in the group must be previously defined in Step 2.</p>
6.	end	Return to privileged EXEC mode.
7.	show running-config	Verify your entries.

Step	Command	Purpose
8.	copy running-config startup-config	(Optional) Save your entries in the configuration file.
9.		Enable RADIUS login authentication. See the “Configuring RADIUS Login Authentication” section on page 137.

To remove the specified RADIUS server, use the **no radius-server host *hostname* | *ip-address*** global configuration command. To remove a server group from the configuration list, use the **no aaa group server radius *group-name*** global configuration command. To remove the IP address of a RADIUS server, use the **no server *ip-address*** server group configuration command.

In this example, the switch is configured to recognize two different RADIUS group servers (*group1* and *group2*). Group1 has two different host entries on the same RADIUS server configured for the same services. The second host entry acts as a fail-over backup to the first entry.

```
Switch(config)# radius-server host 172.20.0.1 auth-port 1000 acct-port 1001
Switch(config)# radius-server host 172.10.0.1 auth-port 1645 acct-port 1646
Switch(config)# aaa new-model
Switch(config)# aaa group server radius group1
Switch(config-sg-radius)# server 172.20.0.1 auth-port 1000 acct-port 1001
Switch(config-sg-radius)# exit
Switch(config)# aaa group server radius group2
Switch(config-sg-radius)# server 172.20.0.1 auth-port 2000 acct-port 2001
Switch(config-sg-radius)# exit
```

Configuring RADIUS Authorization for User Privileged Access and Network Services

AAA authorization limits the services available to a user. When AAA authorization is enabled, the switch uses information retrieved from the user’s profile, which is in the local user database or on the security server, to configure the user’s session. The user is granted access to a requested service only if the information in the user profile allows it.

You can use the **aaa authorization** global configuration command with the **radius** keyword to set parameters that restrict a user’s network access to privileged EXEC mode.

The **aaa authorization exec radius local** command sets these authorization parameters:

- Use RADIUS for privileged EXEC access authorization if authentication was performed by using RADIUS.
- Use the local database if authentication was not performed by using RADIUS.

Note: Authorization is bypassed for authenticated users who log in through the CLI even if authorization has been configured.

Beginning in privileged EXEC mode, follow these steps to specify RADIUS authorization for privileged EXEC access and network services:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	aaa authorization network radius	Configure the switch for user RADIUS authorization for all network-related service requests.

Step	Command	Purpose
3.	aaa authorization exec radius	Configure the switch for user RADIUS authorization to determine if the user has privileged EXEC access. The exec keyword might return user profile information (such as autocommand information).
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable authorization, use the **no aaa authorization {network | exec} method1** global configuration command.

Starting RADIUS Accounting

The AAA accounting feature tracks the services that users are accessing and the amount of network resources that they are consuming. When AAA accounting is enabled, the switch reports user activity to the RADIUS security server in the form of accounting records. Each accounting record contains accounting attribute-value (AV) pairs and is stored on the security server. This data can then be analyzed for network management, client billing, or auditing.

Beginning in privileged EXEC mode, follow these steps to enable RADIUS accounting for each Cisco IOS privilege level and for network services:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	aaa accounting network start-stop radius	Enable RADIUS accounting for all network-related service requests.
3.	aaa accounting exec start-stop radius	Enable RADIUS accounting to send a start-record accounting notice at the beginning of a privileged EXEC process and a stop-record at the end.
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable accounting, use the **no aaa accounting {network | exec} {start-stop} method1...** global configuration command.

Configuring Settings for All RADIUS Servers

Beginning in privileged EXEC mode, follow these steps to configure global communication settings between the switch and all RADIUS servers:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	radius-server key <i>string</i>	Specify the shared secret text string used between the switch and all RADIUS servers. Note: The key is a text string that must match the encryption key used on the RADIUS server. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.
3.	radius-server retransmit <i>retries</i>	Specify the number of times the switch sends each RADIUS request to the server before giving up. The default is 3; the range is 1 to 1000.
4.	radius-server timeout <i>seconds</i>	Specify the number of seconds a switch waits for a reply to a RADIUS request before resending the request. The default is 5 seconds; the range is 1 to 1000.
5.	radius-server deadtime <i>minutes</i>	Specify the number of minutes a RADIUS server, which is not responding to authentication requests, to be skipped, thus avoiding the wait for the request to timeout before trying the next configured server. The default is 0; the range is 1 to 1440 minutes.
6.	end	Return to privileged EXEC mode.
7.	show running-config	Verify your settings.
8.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting for the retransmit, timeout, and deadtime, use the **no** forms of these commands.

Configuring the Switch to Use Vendor-Specific RADIUS Attributes

The Internet Engineering Task Force (IETF) draft standard specifies a method for communicating vendor-specific information between the switch and the RADIUS server by using the vendor-specific attribute (attribute 26). Vendor-specific attributes (VSAs) allow vendors to support their own extended attributes not suitable for general use. The Cisco RADIUS implementation supports one vendor-specific option by using the format recommended in the specification. Cisco's vendor-ID is 9, and the supported option has vendor-type 1, which is named *cisco-avpair*. The value is a string with this format:

```
protocol : attribute sep value *
```

Protocol is a value of the Cisco protocol attribute for a particular type of authorization. *Attribute* and *value* are an appropriate attribute-value (AV) pair defined in the Cisco TACACS+ specification, and *sep* is = for mandatory attributes and is * for optional attributes. The full set of features available for TACACS+ authorization can then be used for RADIUS.

For example, this AV pair activates Cisco's *multiple named ip address pools* feature during IP authorization (during PPP's IPCP address assignment):

```
cisco-avpair= "ip:addr-pool=first"
```

This example shows how to provide a user logging in from a switch with immediate access to privileged EXEC commands:

```
cisco-avpair= "shell:priv-lvl=15"
```

This example shows how to specify an authorized VLAN in the RADIUS server database:

```
cisco-avpair= "tunnel-type(#64)=VLAN(13)"
cisco-avpair= "tunnel-medium-type(#65)=802 media(6)"
cisco-avpair= "tunnel-private-group-ID(#81)=vlanid"
```

Other vendors have their own unique vendor-IDs, options, and associated VSAs. For more information about vendor-IDs and VSAs, refer to RFC 2138, "Remote Authentication Dial-In User Service (RADIUS)."

Beginning in privileged EXEC mode, follow these steps to configure the switch to recognize and use VSAs:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	radius-server vsa send [accounting authentication]	<p>Enable the switch to recognize and use VSAs as defined by RADIUS IETF attribute 26.</p> <ul style="list-style-type: none"> (Optional) Use the accounting keyword to limit the set of recognized vendor-specific attributes to only accounting attributes. (Optional) Use the authentication keyword to limit the set of recognized vendor-specific attributes to only authentication attributes. <p>If you enter this command without keywords, both accounting and authentication vendor-specific attributes are used.</p>
3.	end	Return to privileged EXEC mode.
4.	show running-config	Verify your settings.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

For a complete list of RADIUS attributes or more information about vendor-specific attribute 26, refer to the *Cisco IOS Security Configuration Guide for Cisco IOS Release 12.1*.

Configuring the Switch for Vendor-Proprietary RADIUS Server Communication

Although an IETF draft standard for RADIUS specifies a method for communicating vendor-proprietary information between the switch and the RADIUS server, some vendors have extended the RADIUS attribute set in a unique way. Cisco IOS software supports a subset of vendor-proprietary RADIUS attributes.

As mentioned earlier, to configure RADIUS (whether vendor-proprietary or IETF draft-compliant), you must specify the host running the RADIUS server daemon and the secret text string it shares with the switch. You specify the RADIUS host and secret text string by using the **radius-server** global configuration commands.

Beginning in privileged EXEC mode, follow these steps to specify a vendor-proprietary RADIUS server host and a shared secret text string:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	radius-server host { <i>hostname</i> <i>ip-address</i> } non-standard	Specify the IP address or host name of the remote RADIUS server host and identify that it is using a vendor-proprietary implementation of RADIUS.
3.	radius-server key <i>string</i>	Specify the shared secret text string used between the switch and the vendor-proprietary RADIUS server. The switch and the RADIUS server use this text string to encrypt passwords and exchange responses. Note: The key is a text string that must match the encryption key used on the RADIUS server. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your settings.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete the vendor-proprietary RADIUS host, use the **no radius-server host** {*hostname* | *ip-address*} **non-standard** global configuration command. To disable the key, use the **no radius-server key** global configuration command.

This example shows how to specify a vendor-proprietary RADIUS host and to use a secret key of *rad124* between the switch and the server:

```
Switch(config)# radius-server host 172.20.30.15 nonstandard
Switch(config)# radius-server key rad124
```

Displaying the RADIUS Configuration

To display the RADIUS configuration, use the **show running-config** privileged EXEC command.

Configuring the Switch for Local Authentication and Authorization

You can configure AAA to operate without a server by setting the switch to implement AAA in local mode. The switch then handles authentication and authorization. No accounting is available in this configuration.

Beginning in privileged EXEC mode, follow these steps to configure the switch for local AAA:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	aaa new-model	Enable AAA.

Step	Command	Purpose
3.	aaa authentication login default local	Set the login authentication to use the local username database. The default keyword applies the local user database authentication to all interfaces.
4.	aaa authorization exec local	Configure user AAA authorization to determine if the user is allowed to run an EXEC shell by checking the local database.
5.	aaa authorization network local	Configure user AAA authorization for all network-related service requests.
6.	username <i>name</i> [privilege <i>level</i>] { password <i>encryption-type</i> <i>password</i> }	Enter the local database, and establish a username-based authentication system. Repeat this command for each user. <ul style="list-style-type: none"> For <i>name</i>, specify the user ID as one word. Spaces and quotation marks are not allowed. (Optional) For <i>level</i>, specify the privilege level the user has after gaining access. The range is 0 to 15. Level 15 gives privileged EXEC mode access. Level 0 gives user EXEC mode access. For <i>encryption-type</i>, enter 0 to specify that an unencrypted password follows. Enter 7 to specify that a hidden password follows. For <i>password</i>, specify the password the user must enter to gain access to the switch. The password must be from 1 to 25 characters, can contain embedded spaces, and must be the last option specified in the username command.
7.	end	Return to privileged EXEC mode.
8.	show running-config	Verify your entries.
9.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable AAA, use the **no aaa new-model** global configuration command. To disable authorization, use the **no aaa authorization {network | exec} method1** global configuration command.

Configuring the Switch for Secure Shell

This section describes how to configure the Secure Shell (SSH) feature. SSH is a cryptographic security feature that is subject to export restrictions. To use this feature, the cryptographic (encrypted) software image must be installed on your switch. You must obtain authorization to use this feature and to download the cryptographic software files. For more information, see the “Cryptographic Software Image Guidelines” section. This section contains this configuration information:

- Understanding SSH, on page 147
- Cryptographic Software Image Guidelines, on page 147
- Configuring SSH, on page 147

Note: For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Security Command Reference for Cisco IOS Release 12.2*.

Understanding SSH

SSH is a protocol that provides a secure, remote connection to a device. There are two versions of SSH: SSH version 1 and SSH version 2. This software release only supports SSH version 1.

SSH provides more security for remote connections than Telnet by providing strong encryption when a device is authenticated. The SSH feature has an SSH server and an SSH integrated client. SSH supports these user authentication methods:

- TACACS+ (for more information, see the “Controlling Switch Access with TACACS+” section on page 124)
- RADIUS (for more information, see the “Controlling Switch Access with RADIUS” section on page 131)
- Local authentication and authorization (for more information, see the “Configuring the Switch for Local Authentication and Authorization” section on page 145)

For more information about SSH, refer to the “*Configuring Secure Shell*” section in the *Cisco IOS Security Configuration Guide for Cisco IOS Release 12.2*.

Note: The SSH feature in this software release does not support IP Security (IPSec).

Cryptographic Software Image Guidelines

The SSH feature uses a large amount of switch memory, which limits the number of VLANs, trunk ports, and cluster members that you can configure on the switch. Before you download the cryptographic software image, your switch configuration must meet these conditions:

- The number of trunk ports multiplied by the number of VLANs on the switch must be less than or equal to 256. These are examples of switch configurations that meet this condition:
 - If the switch has 4 trunk ports, it can have up to 64 VLANs.
 - If the switch has 32 VLANs, it can have up to 8 trunk ports.
- If your switch is a cluster command switch, it can only support up to eight cluster members.

To access the cryptographic version of the CIGESM software go to the IBM web site:www.ibm.com/support. Click on the **Download and Driver** icon on the web page. This will bring up a search web page. In the **Search** box type in Cisco and then click submit. This will bring up a web page with a description of the latest software drivers for the CIGESM. Click on this text and you will be directed to a web page listing all the software drivers available for the CIGESM. Find the latest level cryptographic version and then right click on this. A drop-down menu next to the file is displayed. Click on the **Save target as...** to save the file to your hard disk. From your hard disk you can ftp the file to your CIGESM. This process is described in the Command Reference using the **archive download** command.

Configuring SSH

Before configuring SSH, download the cryptographic software image from www.ibm.com/support.

For information about configuring SSH and displaying SSH settings, refer to the “*Configuring Secure Shell*” section in the *Cisco IOS Security Configuration Guide for Cisco IOS Release 12.2*.

Chapter 8. Configuring 802.1X Port-Based Authentication

This chapter describes how to configure IEEE 802.1X port-based authentication on the switch to prevent unauthorized devices (clients) from gaining access to the network.

Note: For complete syntax and usage information for the commands used in this chapter, refer to the *command reference* for this release.

This chapter consists of these sections:

- Understanding 802.1X Port-Based Authentication, on page 149
- Configuring 802.1X Authentication, on page 155
- Displaying 802.1X Statistics and Status, on page 164

Understanding 802.1X Port-Based Authentication

The IEEE 802.1X standard defines a client-server-based access control and authentication protocol that restricts unauthorized clients from connecting to a LAN through publicly accessible ports. The authentication server authenticates each client connected to a switch port before making available any services offered by the switch or the LAN.

Until the client is authenticated, 802.1X access control allows only Extensible Authentication Protocol over LAN (EAPOL), Cisco Discovery Protocol (CDP), and Spanning Tree Protocol (STP) traffic through the port to which the client is connected. After authentication is successful, normal traffic can pass through the port.


These sections describe 802.1X port-based authentication:

- Device Roles, on page 149
- Authentication Initiation and Message Exchange, on page 151
- Ports in Authorized and Unauthorized States, on page 152
- Supported Topologies, on page 152
- Using 802.1X with Port Security, on page 153
- Using 802.1X with VLAN Assignment, on page 154
- Using 802.1X with Guest VLAN, on page 155

Device Roles

With 802.1X port-based authentication, the devices in the network have specific roles as shown in Figure 31.

Figure 31. 802.1X Device Roles



- *Client*—the device (workstation) that requests access to the LAN and switch services and responds to requests from the switch. The workstation must be running 802.1X-compliant client software such as that offered in the Microsoft Windows XP operating system. (The client is the *supplicant* in the IEEE 802.1X specification.)

Note: To resolve Windows XP network connectivity and 802.1X authentication issues, read the Microsoft Knowledge Base article at this URL:
<http://support.microsoft.com/support/kb/articles/Q303/5/97.ASP>

- *Authentication server*—performs the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch whether or not the client is authorized to access the LAN and switch services. Because the switch acts as the proxy, the authentication service is transparent to the client. In this release, the Remote Authentication Dial-In User Service (RADIUS) security system with Extensible Authentication Protocol (EAP) extensions is the only supported authentication server. It is available in Cisco Secure Access Control Server version 3.0 or later. RADIUS operates in a client/server model in which secure authentication information is exchanged between the RADIUS server and one or more RADIUS clients.
- *Switch* (edge switch or wireless access point)—controls the physical access to the network based on the authentication status of the client. The switch acts as an intermediary (proxy) between the client and the authentication server, requesting identity information from the client, verifying that information with the authentication server, and relaying a response to the client. The switch includes the RADIUS client, which is responsible for encapsulating and decapsulating the EAP frames and interacting with the authentication server.

When the switch receives EAPOL frames and relays them to the authentication server, the Ethernet header is stripped and the remaining EAP frame is re-encapsulated in the RADIUS format. The EAP frames are not modified or examined during encapsulation, and the authentication server must support EAP within the native frame format. When the switch receives frames from the authentication server, the server's frame header is removed, leaving the EAP frame, which is then encapsulated for Ethernet and sent to the client.

The devices that can act as intermediaries include the Cisco Systems Intelligent Gigabit Ethernet Switch Modules, Catalyst 3750, Catalyst 3550, Catalyst 2970, Catalyst 2955, Catalyst 2950, Catalyst 2940 switches, or a wireless access point. These devices must be running software that supports the RADIUS client and 802.1X.

Authentication Initiation and Message Exchange

The switch or the client can initiate authentication. If you enable authentication on a port by using the **dot1x port-control auto** interface configuration command, the switch must initiate authentication when it determines that the port link state transitions from down to up. It then sends an EAP-request/identity frame to the client to request its identity (typically, the switch sends an initial identity/request frame followed by one or more requests for authentication information). Upon receipt of the frame, the client responds with an EAP-response/identity frame.

However, if during bootup, the client does not receive an EAP-request/identity frame from the switch, the client can initiate authentication by sending an EAPOL-start frame, which prompts the switch to request the client's identity.


Note: If 802.1X is not enabled or supported on the network access device, any EAPOL frames from the client are dropped. If the client does not receive an EAP-request/identity frame after three attempts to start authentication, the client sends frames as if the port is in the authorized state. A port in the authorized state effectively means that the client has been successfully authenticated. For more information, see the "Ports in Authorized and Unauthorized States" section on page 152.

When the client supplies its identity, the switch begins its role as the intermediary, passing EAP frames between the client and the authentication server until authentication succeeds or fails. If the authentication succeeds, the switch port becomes authorized. For more information, see the "Ports in Authorized and Unauthorized States" section on page 152.

The specific exchange of EAP frames depends on the authentication method being used. Figure 32 shows a message exchange initiated by the client using the One-Time-Password (OTP) authentication method with a RADIUS server.

Note: Authentication of clients connecting through the external 10/100/1000 Mbps ports does not occur unless you disable the default blocking of external ports from the switch.

Figure 32. Message Exchange



Ports in Authorized and Unauthorized States

The switch port state determines whether or not the client is granted access to the network. The port starts in the *unauthorized* state. While in this state, the port disallows all ingress and egress traffic except for 802.1X protocol packets. When a client is successfully authenticated, the port transitions to the *authorized* state, allowing all traffic for the client to flow normally.

If a client that does not support 802.1X is connected to an unauthorized 802.1X port, the switch requests the client's identity. In this situation, the client does not respond to the request, the port remains in the unauthorized state, and the client is not granted access to the network.

In contrast, when an 802.1X-enabled client connects to a port that is not running the 802.1X protocol, the client initiates the authentication process by sending the EAPOL-start frame. When no response is received, the client sends the request for a fixed number of times. Because no response is received, the client begins sending frames as if the port is in the authorized state.

You control the port authorization state by using the **dot1x port-control** interface configuration command and these keywords:

- **force-authorized**—disables 802.1X authentication and causes the port to transition to the authorized state without any authentication exchange required. The port sends and receives normal traffic without 802.1X-based authentication of the client. This is the default setting.
- **force-unauthorized**—causes the port to remain in the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the interface.
- **auto**—enables 802.1X authentication and causes the port to begin in the unauthorized state, allowing only EAPOL frames to be sent and received through the port. The authentication process begins when the link state of the port transitions from down to up or when an EAPOL-start frame is received. The switch requests the identity of the client and begins relaying authentication messages between the client and the authentication server. Each client attempting to access the network is uniquely identified by the switch by using the client's MAC address.

If the client is successfully authenticated (receives an Accept frame from the authentication server), the port state changes to authorized, and all frames from the authenticated client are allowed through the port. If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the switch can resend the request. If no response is received from the server after the specified number of attempts, authentication fails, and network access is not granted.

When a client logs off, it sends an EAPOL-logoff message, causing the switch port to transition to the unauthorized state.

If the link state of a port transitions from up to down, or if an EAPOL-logoff frame is received, the port returns to the unauthorized state.

Supported Topologies


The 802.1X port-based authentication is supported in two topologies:

- Point-to-point
- Wireless LAN

In a point-to-point configuration (see Figure 31. on page 150), only one client can be connected to the 802.1X-enabled switch port. The switch detects the client when the port link state changes to the up state. If a client leaves or is replaced with another client, the switch changes the port link state to down, and the port returns to the unauthorized state.

Figure 33 shows 802.1X port-based authentication in a wireless LAN. The 802.1X port is configured as a multiple-hosts port that becomes authorized as soon as one client is authenticated. When the port is authorized, all other hosts indirectly attached to the port are granted access to the network. If the port becomes unauthorized (re-authentication fails or an EAPOL-logoff message is received), the switch denies access to the network to all of the attached clients. In this topology, the wireless access point is responsible for authenticating the clients attached to it, and the wireless access point acts as a client to the switch.

Figure 33. Wireless LAN Example



Using 802.1X with Port Security

You can enable an 802.1X port for port security in either single-host or multiple-hosts mode. (You must also configure port security on the port by using the **switchport port-security** interface configuration command.) When you enable port security and 802.1X on a port, 802.1X authenticates the port, and port security manages network access for all MAC addresses, including that of the client. You can then limit the number or group of clients that can access the network through an 802.1X port.

These are some examples of the interaction between 802.1X and port security on the switch:

- When a client is authenticated, and the port security table is not full, the client's MAC address is added to the port security list of secure hosts. The port then proceeds to come up normally.

When a client is authenticated and manually configured for port security, it is guaranteed an entry in the secure host table (unless port security static aging has been enabled).

A security violation occurs if the client is authenticated, but port security table is full. This can happen if the maximum number of secure hosts has been statically configured, or if the client ages out of the secure host table. If the client's address is aged out, its place in the secure host table can be taken by another host.

The port security violation modes determine the action for security violations. For more information, see the “Security Violations” section on page 319.

- When an 802.1X client logs off, the port transitions back to an unauthenticated state, and all dynamic entries in the secure host table are cleared, including the entry for the client. Normal authentication then takes place.
- If the port is administratively shut down, the port becomes unauthenticated, and all dynamic entries are removed from the secure host table.
- Port security and a voice VLAN can be configured simultaneously on an 802.1X port that is in either single-host or multiple-hosts mode. Port security applies to both the voice VLAN identifier (VVID) and the port VLAN identifier (PVID).
- When an 802.1X client address is manually removed from the port security table, we recommend that you re-authenticate the client by entering the **dot1x re-authenticate** privileged EXEC command.

For more information about enabling port security on your switch, see the “Configuring Port Security” section on page 318.

Using 802.1X with VLAN Assignment

You can limit network access for certain users by using VLAN assignment. After successful 802.1X authentication of a port, the RADIUS server sends the VLAN assignment to configure the switch port. The RADIUS server database maintains the username-to-VLAN mappings, which assigns the VLAN based on the username of the client connected to the switch port.

When configured on the switch and the RADIUS server, 802.1X with VLAN assignment has these characteristics:

- If no VLAN is supplied by the RADIUS server or if 802.1X authorization is disabled, the port is configured in its access VLAN after successful authentication.
- If 802.1X authorization is enabled but the VLAN information from the RADIUS server is not valid, the port returns to the unauthorized state and remains in the configured access VLAN. This prevents ports from appearing unexpectedly in an inappropriate VLAN because of a configuration error.

Configuration errors could include a VLAN specified for a routed port, a malformed VLAN ID, a nonexistent or internal (routed port) VLAN ID, or attempted assignment to a voice VLAN ID.

- If 802.1X authorization is enabled and all information from the RADIUS server is valid, the port is placed in the specified VLAN after authentication.
- If the multiple-hosts mode is enabled on an 802.1X port, all hosts are placed in the same VLAN (specified by the RADIUS server) as the first authenticated host.
- If port security is enabled on an 802.1X port with VLAN assignment, the port is placed in the RADIUS server assigned VLAN.
- If 802.1X is disabled on the port, it is returned to the configured access VLAN.
- When the port is in the force authorized, force unauthorized, unauthorized, or shutdown state, it is placed in the configured access VLAN.
- If an 802.1X port is authenticated and put in the RADIUS server assigned VLAN, any change to the port access VLAN configuration does not take effect.
- The 802.1X with VLAN assignment feature is not supported on trunk ports, dynamic ports, or with dynamic-access port assignment through a VLAN Membership Policy Server (VMPS).

To configure VLAN assignment you need to perform these tasks:

- Enable AAA authorization.
- Enable 802.1X (the VLAN assignment feature is automatically enabled when you configure 802.1X on an access port).
- Assign vendor-specific tunnel attributes in the RADIUS server. The RADIUS server must return these attributes to the switch:
 - [64] Tunnel-Type = VLAN
 - [65] Tunnel-Medium-Type = 802
 - [81] Tunnel-Private-Group-ID = VLAN name or VLAN ID
 Attribute [64] must contain the value *VLAN* (type 13). Attribute [65] must contain the value *802* (type 6). Attribute [81] specifies the *VLAN name* or *VLAN ID* assigned to the 802.1X-authenticated user.

For examples of tunnel attributes, see the “Configuring the Switch to Use Vendor-Specific RADIUS Attributes” section on page 143.

Using 802.1X with Guest VLAN

You can configure a guest VLAN for each 802.1X port on the switch to provide limited services to clients (for example, how to download the 802.1X client). These clients might be upgrading their system for 802.1X authentication, and some hosts, such as Windows 98 systems, might not be 802.1X-capable.

When the authentication server does not receive a response to its EAPOL request/identity frame, clients that are not 802.1X-capable are put into the guest VLAN for the port, if one is configured. However, the server does not grant 802.1X-capable clients that fail authentication access to the network. Any number of hosts are allowed access once the switch port is moved to the guest VLAN. If an 802.1X-capable host joins the same port on which the guest VLAN is configured, the port is put into the unauthorized state in the user-configured access VLAN, and authentication is restarted.

Guest VLANs are supported on 802.1X ports in single-host and multiple-hosts modes.

You can configure any VLAN, except RSPAN VLANs or voice VLAN IDs (VVIDs), as an 802.1X guest VLAN. The guest VLAN feature is not supported on trunk ports; it is supported only on access ports.

For configuration steps, see the “Configuring a Guest VLAN” section on page 163.

Configuring 802.1X Authentication

These sections describe how to configure 802.1X port-based authentication on your switch:

- Default 802.1X Configuration, on page 156
- 802.1X Configuration Guidelines, on page 156
- Enabling 802.1X Authentication, on page 157 (required)
- Configuring the Switch-to-RADIUS-Server Communication, on page 159 (required)
- Enabling Periodic Re-Authentication, on page 160 (optional)
- Manually Re-Authenticating a Client Connected to a Port, on page 160 (optional)
- Changing the Quiet Period, on page 161 (optional)
- Changing the Switch-to-Client Retransmission Time, on page 161 (optional)

- Setting the Switch-to-Client Frame-Retransmission Number, on page 162 (optional)
- Configuring the Host Mode, on page 162 (optional)
- Configuring a Guest VLAN, on page 163 (optional)
- Resetting the 802.1X Configuration to the Default Values, on page 164 (optional)

Default 802.1X Configuration

Table 32 shows the default 802.1X configuration.

Table 32. Default 802.1X Configuration

Feature	Default Setting
Authentication, authorization, and accounting (AAA)	Disabled
RADIUS server	
<ul style="list-style-type: none"> • IP address • UDP authentication port • Key 	<ul style="list-style-type: none"> • None specified • 1812 • None specified
Switch 802.1X enable state	Disabled
Per-interface 802.1X enable state	Disabled (force-authorized) The port sends and receives normal traffic without 802.1X-based authentication of the client.
Periodic re-authentication	Disabled
Number of seconds between re-authentication attempts	3600 seconds
Quiet period	60 seconds (number of seconds that the switch remains in the quiet state following a failed authentication exchange with the client)
Retransmission time	30 seconds (number of seconds that the switch should wait for a response to an EAP request/identity frame from the client before resending the request)
Maximum retransmission number	2 times (number of times that the switch will send an EAP-request/identity frame before restarting the authentication process)
Host mode	Single-host mode
Guest VLAN	None specified
Client timeout period	30 seconds (when relaying a request from the authentication server to the client, the amount of time the switch waits for a response before resending the request to the client)
Authentication server timeout period	30 seconds (when relaying a response from the client to the authentication server, the amount of time the switch waits for a reply before resending the response to the server. This setting is not configurable)

802.1X Configuration Guidelines

These are the 802.1X authentication configuration guidelines:

- When 802.1X is enabled, ports are authenticated before any other Layer 2 features are enabled.
- The 802.1X protocol is supported on Layer 2 static-access ports and voice VLAN ports, but it is not supported on these port types:
 - Trunk port—If you try to enable 802.1X on a trunk port, an error message appears, and 802.1X is not enabled. If you try to change the mode of an 802.1X-enabled port to trunk, the port mode is not changed.
 - Dynamic ports—A port in dynamic mode can negotiate with its neighbor to become a trunk port. If you try to enable 802.1X on a dynamic port, an error message appears, and 802.1X is not enabled. If you try to change the mode of an 802.1X-enabled port to dynamic, the port mode is not changed.
 - Dynamic-access ports—If you try to enable 802.1X on a dynamic-access (VLAN Query Protocol [VQP]) port, an error message appears, and 802.1X is not enabled. If you try to change an 802.1X-enabled port to dynamic VLAN assignment, an error message appears, and the VLAN configuration is not changed.
 - EtherChannel port—Before enabling 802.1X on the port, you must first remove it from the EtherChannel. If you try to enable 802.1X on an EtherChannel or on an active port in an EtherChannel, an error message appears, and 802.1X is not enabled. If you enable 802.1X on a not-yet active port of an EtherChannel, the port does not join the EtherChannel.
 - Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) destination ports—You can enable 802.1X on a port that is a SPAN or RSPAN destination or reflector port. However, 802.1X is disabled until the port is removed as a SPAN or RSPAN destination or reflector port. You can enable 802.1X on a SPAN or RSPAN source port.
- You can configure any VLAN, except RSPAN VLANs or voice VVIDs, as an 802.1X guest VLAN. The guest VLAN feature is not supported trunk ports; it is supported only on access ports.
- When 802.1X is enabled on a port, you cannot configure a port VLAN that is equal to a voice VLAN.
- The 802.1X with VLAN assignment feature is not supported on trunk ports, dynamic ports, or with dynamic-access port assignment through a VMPS.

Enabling 802.1X Authentication

To enable 802.1X port-based authentication, you must enable AAA and specify the authentication method list. A method list describes the sequence and authentication methods to be queried to authenticate a user.

The software uses the first method listed to authenticate users; if that method fails to respond, the software selects the next authentication method in the method list. This process continues until there is successful communication with a listed authentication method or until all defined methods are exhausted. If authentication fails at any point in this cycle, the authentication process stops, and no other authentication methods are attempted.

To allow VLAN assignment, you must enable AAA authorization to configure the switch for all network-related service requests.

Beginning in privileged EXEC mode, follow these steps to configure 802.1X port-based authentication. This procedure is required.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	aaa new-model	Enable AAA.
3.	aaa authentication dot1x {default} method1 [method2...]	Create an 802.1X authentication method list. To create a default list that is used when a named list is <i>not</i> specified in the authentication command, use the default keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all interfaces. Enter at least one of these keywords: <ul style="list-style-type: none"> • group radius—Use the list of all RADIUS servers for authentication. • none—Use no authentication. The client is automatically authenticated by the switch without using the information supplied by the client.
4.	dot1x system-auth-control	Enable 802.1X authentication globally on the switch.
5.	aaa authorization network {default} group radius	(Optional) Configure the switch for user RADIUS authorization for all network-related service requests, such as VLAN assignment. Note:
6.	interface interface-id	Enter interface configuration mode, and specify the interface connected to the client to be enabled for 802.1X authentication.
7.	dot1x port-control auto	Enable 802.1X authentication on the interface. For feature interaction information, see the “802.1X Configuration Guidelines” section on page 156.
8.	end	Return to privileged EXEC mode.
9.	show dot1x	Verify your entries. Check the Status column in the 802.1X Port Summary section of the display. An <i>enabled</i> status means the port-control value is set either to auto or to force-unauthorized .
10.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable AAA, use the **no aaa new-model** global configuration command. To disable 802.1X AAA authentication, use the **no aaa authentication dot1x {default | list-name}** global configuration command. To disable 802.1X AAA authorization, use the **no aaa authorization** global configuration command. To disable 802.1X authentication on the switch, use the **no dot1x system-auth-control** global configuration command.

This example shows how to enable AAA and 802.1X on Gigabit Ethernet port 17:

```
Switch# configure terminal
Switch(config)# aaa new-model
Switch(config)# aaa authentication dot1x default group radius
Switch(config)# dot1x system-auth-control
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# switchport mode access
Switch(config-if)# dot1x port-control auto
```



```
Switch(config-if)# end
```

Configuring the Switch-to-RADIUS-Server Communication

RADIUS security servers are identified by their host name or IP address, host name and specific UDP port numbers, or IP address and specific UDP port numbers. The combination of the IP address and UDP port number creates a unique identifier, which enables RADIUS requests to be sent to multiple UDP ports on a server at the same IP address. If two different host entries on the same RADIUS server are configured for the same service—for example, authentication—the second host entry configured acts as the fail-over backup to the first one. The RADIUS host entries are tried in the order that they were configured.

Beginning in privileged EXEC mode, follow these steps to configure the RADIUS server parameters on the switch. This procedure is required.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	radius-server host { <i>hostname</i> <i>ip-address</i> } auth-port <i>port-number</i> key <i>string</i>	Configure the RADIUS server parameters on the switch. For <i>hostname</i> <i>ip-address</i> , specify the host name or IP address of the remote RADIUS server. For auth-port <i>port-number</i> , specify the UDP destination port for authentication requests. The default is 1812. For key <i>string</i> , specify the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. The key is a text string that must match the encryption key used on the RADIUS server. Note: Always configure the key as the last item in the radius-server host command syntax because leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in the key, do not enclose the key in quotation marks unless the quotation marks are part of the key. This key must match the encryption used on the RADIUS daemon. If you want to use multiple RADIUS servers, re-enter this command.
3.	end	Return to privileged EXEC mode.
4.	show running-config	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete the specified RADIUS server, use the **no radius-server host** {*hostname* | *ip-address*} global configuration command.

This example shows how to specify the server with IP address 172.20.39.46 as the RADIUS server, to use port 1612 as the authorization port, and to set the encryption key to *rad123*, matching the key on the RADIUS server:

```
Switch(config)# radius-server host 172.120.39.46 auth-port 1612 key rad123
```

You can globally configure the timeout, retransmission, and encryption key values for all RADIUS servers by using the **radius-server host** global configuration command. If you want to configure these options on a per-server basis, use the **radius-server**

timeout, **radius-server retransmit**, and the **radius-server key** global configuration commands. For more information, see the “Configuring Settings for All RADIUS Servers” section on page 142.

You also need to configure some settings on the RADIUS server. These settings include the IP address of the switch and the key string to be shared by both the server and the switch. For more information, refer to the RADIUS server documentation.

Enabling Periodic Re-Authentication

You can enable periodic 802.1X client re-authentication and specify how often it occurs. If you do not specify a time period before enabling re-authentication, the number of seconds between re-authentication attempts is 3600.

Beginning in privileged EXEC mode, follow these steps to enable periodic re-authentication of the client and to configure the number of seconds between re-authentication attempts. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be configured.
3.	dot1x reauthentication	Enable periodic re-authentication of the client, which is disabled by default.
4.	dot1x timeout reauth-period <i>seconds</i>	Set the number of seconds between re-authentication attempts. The range is 1 to 65535; the default is 3600 seconds. This command affects the behavior of the switch only if periodic re-authentication is enabled.
5.	end	Return to privileged EXEC mode.
6.	show dot1x interface <i>interface-id</i>	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable periodic re-authentication, use the **no dot1x reauthentication** interface configuration command. To return to the default number of seconds between re-authentication attempts, use the **no dot1x timeout reauth-period** global configuration command.

This example shows how to enable periodic re-authentication and set the number of seconds between re-authentication attempts to 4000:

```
Switch(config-if)# dot1x reauthentication
Switch(config-if)# dot1x timeout reauth-period 4000
```

Manually Re-Authenticating a Client Connected to a Port

You can manually re-authenticate the client connected to a specific port at any time by entering the **dot1x re-authenticate interface** *interface-id* privileged EXEC command. This step is optional. If you want to enable or disable periodic re-authentication, see the “Enabling Periodic Re-Authentication” section on page 160.

This example shows how to manually re-authenticate the client connected to Gigabit Ethernet port 17:

Switch# **dot1x re-authenticate interface gigabitethernet0/17**

Changing the Quiet Period

When the switch cannot authenticate the client, the switch remains idle for a set period of time, and then tries again. The idle time is determined by the quiet-period value. A failed authentication of the client might occur because the client provided an invalid password. You can provide a faster response time to the user by entering a smaller number than the default.

Beginning in privileged EXEC mode, follow these steps to change the quiet period. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be configured.
3.	dot1x timeout quiet-period <i>seconds</i>	Set the number of seconds that the switch remains in the quiet state following a failed authentication exchange with the client. The range is 1 to 65535 seconds; the default is 60.
4.	end	Return to privileged EXEC mode.
5.	show dot1x interface <i>interface-id</i>	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default quiet time, use the **no dot1x timeout quiet-period** interface configuration command.

This example shows how to set the quiet time on the switch to 30 seconds:

```
Switch(config-if)# dot1x timeout quiet-period 30
```

Changing the Switch-to-Client Retransmission Time

The client responds to the EAP-request/identity frame from the switch with an EAP-response/identity frame. If the switch does not receive this response, it waits a set period of time (known as the retransmission time) and then resends the frame.

Note: You should change the default value of this command only to adjust for unusual circumstances such as unreliable links or specific behavioral problems with certain clients and authentication servers.

Beginning in privileged EXEC mode, follow these steps to change the amount of time that the switch waits for client notification. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be configured.
3.	dot1x timeout tx-period <i>seconds</i>	Set the number of seconds that the switch waits for a response to an EAP-request/identity frame from the client before resending the request. The range is 1 to 65535 seconds; the default is 30.
4.	end	Return to privileged EXEC mode.

Step	Command	Purpose
5.	show dot1x interface <i>interface-id</i>	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default retransmission time, use the **no dot1x timeout tx-period** interface configuration command.

This example shows how to set 60 as the number of seconds that the switch waits for a response to an EAP-request/identity frame from the client before resending the request:

```
Switch(config-if)# dot1x timeout tx-period 60
```

Setting the Switch-to-Client Frame-Retransmission Number

In addition to changing the switch-to-client retransmission time, you can change the number of times that the switch sends an EAP-request/identity frame (assuming no response is received) to the client before restarting the authentication process.

Note: You should change the default value of this command only to adjust for unusual circumstances such as unreliable links or specific behavioral problems with certain clients and authentication servers.

Beginning in privileged EXEC mode, follow these steps to set the switch-to-client frame-retransmission number. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be configured.
3.	dot1x max-req <i>count</i>	Set the number of times that the switch sends an EAP-request/identity frame to the client before restarting the authentication process. The range is 1 to 10; the default is 2.
4.	end	Return to privileged EXEC mode.
5.	show dot1x interface <i>interface-id</i>	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default retransmission number, use the **no dot1x max-req** interface configuration command.

This example shows how to set 5 as the number of times that the switch sends an EAP-request/identity request before restarting the authentication process:

```
Switch(config-if)# dot1x max-req 5
```

Configuring the Host Mode

You can configure an 802.1X port for single-host or for multiple-hosts mode. In single-host mode, only one host is allowed on an 802.1X port. When the host is authenticated, the port is placed in the authorized state. When the host leaves the port, the port becomes unauthorized. Packets from hosts other than the authenticated one are dropped.

You can attach multiple hosts to a single 802.1X-enabled port as shown in Figure 33. on page 153. In this mode, only one of the attached hosts must be successfully authorized for all hosts to be granted network access. If the port becomes unauthorized (re-authentication fails or an EAPOL-logoff message is received), all attached clients are denied access to the network.

With the multiple-hosts mode enabled, you can use 802.1X to authenticate the port and port security to manage network access for all MAC addresses, including that of the client.

Beginning in privileged EXEC mode, follow these steps to allow multiple hosts (clients) on an 802.1X-authorized port that has the **dot1x port-control** interface configuration command set to **auto**. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to which multiple hosts are indirectly attached.
3.	dot1x host-mode multi-host	Allow multiple hosts (clients) on an 802.1X-authorized port. Make sure that the dot1x port-control interface configuration command set is set to auto for the specified interface.
4.	end	Return to privileged EXEC mode.
5.	show dot1x interface <i>interface-id</i>	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable multiple hosts on the port, use the **no dot1x host-mode multi-host** interface configuration command.

This example shows how to enable Gigabit Ethernet interface 17 to allow multiple hosts:

```
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# dot1x port-control auto
Switch(config-if)# dot1x host-mode multi-host
```

Configuring a Guest VLAN

When you configure a guest VLAN, clients that are not 802.1X-capable are put into the guest VLAN when the server does not receive a response to its EAPOL request/identity frame. Clients that are 802.1X-capable but fail authentication are not granted access to the network. The switch supports guest VLANs in single-host or multiple-hosts mode.

Beginning in privileged EXEC mode, follow these steps to configure a guest VLAN. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be configured. For the supported interface types, see the “802.1X Configuration Guidelines” section on page 156.

Step	Command	Purpose
3.	dot1x guest-vlan <i>vlan-id</i>	Specify an active VLAN as an 802.1X guest VLAN. The range is 1 to 4094. Any VLAN can be configured as an 802.1X guest VLAN except RSPAN VLANs or voice VLANs.
4.	end	Return to privileged EXEC mode.
5.	show dot1x interface <i>interface-id</i>	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable and remove the guest VLAN, use the **no dot1x guest-vlan** interface configuration command. The port returns to the unauthorized state.

This example shows how to enable VLAN 9 as an 802.1X guest VLAN on Gigabit Ethernet interface 17:

```
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# dot1x guest-vlan 9
```

Resetting the 802.1X Configuration to the Default Values

Beginning in privileged EXEC mode, follow these steps to reset the 802.1X configuration to the default values.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be configured.
3.	dot1x default	Reset the configurable 802.1X parameters to the default values.
4.	end	Return to privileged EXEC mode.
5.	show dot1x interface <i>interface-id</i>	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Displaying 802.1X Statistics and Status

To display 802.1X statistics for all interfaces, use the **show dot1x all statistics** privileged EXEC command. To display 802.1X statistics for a specific interface, use the **show dot1x statistics interface** *interface-id* privileged EXEC command.

To display the 802.1X administrative and operational status for the switch, use the **show dot1x all** privileged EXEC command. To display the 802.1X administrative and operational status for a specific interface, use the **show dot1x interface** *interface-id* privileged EXEC command.

For detailed information about the fields in these displays, refer to the *command reference* for this release.

Chapter 9. Configuring the Switch Interfaces

This chapter describes the types of interfaces on a switch and how to configure them. The chapter has these sections:

- Understanding Interface Types, on page 165
- Using the Interface Command, on page 168
- Configuring Ethernet Interfaces, on page 173
- Monitoring and Maintaining the Interfaces, on page 178

Note: For complete syntax and usage information for the commands used in this chapter, refer to the switch *command reference* for this release and the online *Cisco IOS Interface Command Reference for Cisco IOS Release 12.1*.

Understanding Interface Types

This section describes the different types of interfaces supported by the switch with references to chapters that contain more detailed information about configuring these interface types. The rest of the chapter describes configuration procedures for switch ports.

Note: The switch has 4 external ports and 16 internal ports, as described in the “Specifying Ports in Interface Configuration Mode” section on page 27.

Switch ports are Layer 2-only interfaces associated with a physical port. They are used for managing the physical interface and associated Layer 2 protocols and do not handle routing or bridging. A switch port can be an access port or a trunk port.

You can configure a port as an access port or trunk port or let the Dynamic Trunking Protocol (DTP) operate on a per-port basis to determine if a switch port should be an access port or a trunk port by negotiating with the port on the other end of the link.

Configure switch ports by using the **switchport** interface configuration commands. For detailed information about configuring access port and trunk port characteristics, see Chapter 13 “Configuring VLANs.”

Note: The physical external switch ports are 10/100/1000 Mbps Ethernet ports. For more information, refer to the switch hardware installation guide.

These sections describes these types of interfaces:

- Access Ports, on page 165
- Trunk Ports, on page 166
- Port-Based VLANs, on page 166
- EtherChannel Port Groups, on page 167
- Connecting Interfaces, on page 167

Access Ports

An access port belongs to and carries the traffic of only one VLAN (unless it is configured as a voice VLAN port). Traffic is received and sent in native formats with no VLAN tagging. Traffic arriving on an access port is assumed to belong to the VLAN assigned to the port. If an access port receives an 802.1p- or 802.1Q-tagged packet for the VLAN assigned to the port, the packet is forwarded. If the port receives an 802.1p- or 802.1Q-tagged packet for another VLAN, the packet is dropped, the source address is not learned, and the frame is counted in the *No destination* statistic.

The switch does not support ISL-tagged packets. If the switch receives an ISL-tagged packet, the packet is flooded in the native VLAN of the port on which it was received because the MAC destination address in the ISL-tagged packet is a multicast address.

Two types of access ports are supported:

- Static access ports are manually assigned to a VLAN.
- VLAN membership of dynamic access ports is learned through incoming packets. By default, a dynamic access port is a member of no VLAN, and forwarding to and from the port is enabled only when the VLAN membership of the port is discovered. Dynamic access ports on the switch are assigned to a VLAN by a VLAN Membership Policy Server (VMPS). The switch does not support the function of a VMPS. However, the VMPS can be a Catalyst 6000 series switch.

Trunk Ports

A trunk port carries the traffic of multiple VLANs and by default is a member of all VLANs in the VLAN database. Only IEEE 802.1Q trunk ports are supported. An IEEE 802.1Q trunk port supports simultaneous tagged and untagged traffic. An 802.1Q trunk port is assigned a default Port VLAN ID (PVID), and all untagged traffic travels on the port default PVID. All untagged traffic and tagged traffic with a NULL VLAN ID are assumed to belong to the port default PVID. A packet with a VLAN ID equal to the outgoing port default PVID is sent untagged. All other traffic is sent with a VLAN tag.

Although by default, a trunk port is a member of every VLAN known to the VTP, you can limit VLAN membership by configuring an allowed list of VLANs for each trunk port. The list of allowed VLANs does not affect any other port but the associated trunk port.

By default, the allowed list is different for the switch internal and external ports:

- VLAN ID range is 2 to 4094 on the internal 1000 Mbps ports
- VLAN ID range is 1 on the internal 100 Mbps management module ports
- VLAN ID range is 1 to 4094 on the external 10/100/1000 Mbps ports

A trunk port can only become a member of a VLAN if VTP knows of the VLAN and the VLAN is in the enabled state. If VTP learns of a new, enabled VLAN and the VLAN is in the allowed list for a trunk port, the trunk port automatically becomes a member of that VLAN and traffic is forwarded to and from the trunk port for that VLAN. If VTP learns of a new, enabled VLAN that is not in the allowed list for a trunk port, the port does not become a member of the VLAN, and no traffic for the VLAN is forwarded to or from the port.

For more information about trunk ports, see Chapter 13 “Configuring VLANs.”

Port-Based VLANs

A VLAN is a switched network that is logically segmented by function, team, or application, without regard to the physical location of the users. For more information about VLANs, see Chapter 13 “Configuring VLANs.” Packets received on a port are forwarded only to ports that belong to the same VLAN as the receiving port. Network devices in different VLANs cannot communicate with one another without a Layer 3 device to route traffic between the VLANs.

VLAN partitions provide hard firewalls for traffic in the VLAN, and each VLAN has its own MAC address table. A VLAN comes into existence when a local port is configured to be associated with the VLAN, when the VLAN Trunking Protocol (VTP) learns of its existence from a neighbor on a trunk, or when a user creates a VLAN.

To configure normal-range VLANs (VLAN IDs 1 to 1005), use the **vlan *vlan-id*** global configuration command to enter config-vlan mode or the **vlan database** privileged EXEC command to enter VLAN configuration mode. The VLAN configurations for VLAN IDs 1 to 1005 are saved in the VLAN database. To configure extended-range VLANs (VLAN IDs 1006 to 4094), you must use config-vlan mode with VTP mode set to transparent. Extended-range VLANs are not added to the VLAN database. When VTP mode is transparent, the VTP and VLAN configuration is saved in the switch running configuration, and you can save it in the switch startup configuration file by entering the **copy running-config startup-config** privileged EXEC command.

Add ports to a VLAN by using the **switchport** interface configuration commands:

- Identify the interface.
- For a trunk port, set trunk characteristics, and if desired, define the VLANs to which it can belong.
- For an access port, set and define the VLAN to which it belongs.

EtherChannel Port Groups

EtherChannel port groups provide the ability to treat multiple switch ports as one switch port. These port groups act as a single logical port for high-bandwidth connections between switches or between switches and servers. An EtherChannel balances the traffic load across the links in the channel. If a link within the EtherChannel fails, traffic previously carried over the failed link changes to the remaining links. You can group multiple trunk ports into one logical trunk port or group multiple access ports into one logical access port. Most protocols operate over either single ports or aggregated switch ports and do not recognize the physical ports within the port group. Exceptions are the DTP, the Cisco Discovery Protocol (CDP), the Port Aggregation Protocol (PAgP), and Link Aggregation Control Protocol (LACP) which operate only on physical ports.


When you configure an EtherChannel, you create a port-channel logical interface and assign an interface to the EtherChannel. For Layer 2 interfaces, the logical interface is dynamically created. You manually assign an interface to the EtherChannel by using the **channel-group** interface configuration command. This command binds the physical and logical ports together. For more information, see Chapter 25 “Configuring EtherChannels.”

Connecting Interfaces

Devices within a single VLAN can communicate directly through any switch. Ports in different VLANs cannot exchange data without going through a routing device or routed interface.

With a standard Layer 2 switch, ports in different VLANs have to exchange information through a router. In the configuration shown in Figure 34, when Host A in VLAN 20 sends data to Host B in VLAN 30, it must go from Host A to the switch, to the router, back to the switch, and then to Host B.

Figure 34. Connecting VLANs with Layer 2 Switches



Using the Interface Command

To configure a physical interface (port), use the **interface** global configuration command to enter interface configuration mode and to specify the interface type, slot, and number.

- Type—Gigabit Ethernet (gigabitethernet or gi)
- Slot—The slot number on the switch (always 0 on this switch).
- Port number—The interface number on the switch. The port numbers always begin at 1, starting at the left when facing the front of the switch, for example, gigabitethernet 0/1, gigabitethernet 0/2.

You can identify physical interfaces by physically checking the interface location on the switch. You can also use the **show** privileged EXEC commands to display information about a specific interface or all the interfaces on the switch. The remainder of this chapter primarily provides physical interface configuration procedures.

This section describes how to configure all types of interfaces and how to configure a range of interfaces:

- Procedures for Configuring Interfaces, on page 168
- Configuring a Range of Interfaces, on page 169
- Configuring and Using Interface-Range Macros, on page 171

Procedures for Configuring Interfaces

These general instructions apply to all interface configuration processes.

1. Enter the **configure terminal** command at the privileged EXEC prompt:

```
Switch# configure terminal  
Enter configuration commands, one per line. End with CNTL/Z.  
Switch(config)#
```

2. Enter the **interface** global configuration command. Identify the interface type and the number of the connector. In this example, Gigabit Ethernet interface 0/17 is selected:

```
Switch(config)# interface gigabitethernet0/17
```

```
Switch(config-if)#
```

Note: You do not need to add a space between the interface type and interface number. For example, in the preceding line, you can specify either **gigabitethernet 0/1**, **gigabitethernet0/1**, **gi 0/1**, or **gi0/1**.

3. Follow each **interface** command with the interface configuration commands your particular interface requires. The commands you enter define the protocols and applications that will run on the interface. The commands are collected and applied to the interface when you enter another interface command or enter **end** to return to privileged EXEC mode.

You can also configure a range of interfaces by using the **interface range** or **interface range macro** global configuration commands. Interfaces configured in a range must be the same type and must be configured with the same feature options.

4. After you configure an interface, verify its status by using the **show** privileged EXEC commands listed in the “Monitoring and Maintaining the Interfaces” section on page 178.

Enter the **show interfaces** privileged EXEC command to see a list of all interfaces on or configured for the switch. A report is provided for each interface that the device supports or for the specified interface:

```
Switch# show interfaces
Vlan1 is up, line protocol is up
  Hardware is CPU Interface, address is 000d.ed46.bf00 (bia
000d.ed46.bf00)
  Internet address is 172.20.138.185/28
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 1000 bits/sec, 1 packets/sec
  5 minute output rate 1000 bits/sec, 2 packets/sec
    58834 packets input, 19008862 bytes, 0 no buffer
    Received 435 broadcasts (0 IP multicast)
    0 runts, 0 giants, 0 throttles
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
    1611241 packets output, 117703876 bytes, 0 underruns
    0 output errors, 2 interface resets
    0 output buffer failures, 0 output buffers swapped out
```

<output truncated>

Configuring a Range of Interfaces

You can use the **interface range** global configuration command to configure multiple interfaces with the same configuration parameters. When you enter the interface-range configuration mode, all command parameters that you enter are attributed to all interfaces within that range until you exit this mode.

Beginning in privileged EXEC mode, follow these steps to configure a range of interfaces with the same parameters:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface range { <i>port-range</i> macro <i>macro_name</i> }	Enter interface-range configuration mode by entering the range of interfaces (VLANs or physical ports) to be configured. <ul style="list-style-type: none"> You can use the interface range command to configure up to five port ranges or a previously defined macro. The macro variable is explained in the “Configuring and Using Interface-Range Macros” section on page 171. Each comma-separated <i>port-range</i> must consist of the same port type. You do not need to enter spaces before or after the comma. When you define a range, the space between the first port and the hyphen is required.
3.		You can now use the normal configuration commands to apply the configuration parameters to all interfaces in the range.
4.	end	Return to privileged EXEC mode.
5.	show interfaces [<i>interface-id</i>]	Verify the configuration of the interfaces in the range.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

When using the **interface range** global configuration command, note these guidelines:

- Valid entries for *port-range*:
 - vlan** *vlan-ID* - *vlan-ID*, where VLAN ID is from 1 to 4094
 - gigabitethernet** *slot*{*first port*} - {*last port*}, where slot is **0**
 - port-channel** *port-channel-number* - *port-channel-number*, where *port-channel-number* is from 1 to 6
- You must add a space between the interface numbers and the hyphen when using the **interface range** command. For example, the command **interface range gigabitethernet 0/1 - 5** is a valid range; the command **interface range gigabitethernet 0/1-5** is not a valid range.
- The **interface range** command works only with VLAN interfaces that have been configured with the **interface vlan** command (the **show running-config** privileged EXEC command output shows the configured VLAN interfaces). VLAN interfaces that do not appear by using the **show running-config** command cannot be used with the **interface range** command.
- All interfaces in a range must be the same type; that is, all Gigabit Ethernet ports, all EtherChannel ports, or VLAN interfaces.

This example shows how to use the **interface range** global configuration command to enable Gigabit Ethernet interfaces 0/17 to 0/20:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet0/17 - 20
Switch(config-if-range)# no shutdown
Switch(config-if-range)#
*Oct 6 08:24:35: %LINK-3-UPDOWN: Interface gigabitethernet0/17, changed
state to up
```

```
*Oct 6 08:24:35: %LINK-3-UPDOWN: Interface gigabitethernet0/18, changed
state to up
*Oct 6 08:24:35: %LINK-3-UPDOWN: Interface gigabitethernet0/19, changed
state to up
*Oct 6 08:24:35: %LINK-3-UPDOWN: Interface gigabitethernet0/20, changed
state to up
```

This example shows how to use a comma to add different interface type strings to the range to enable Gigabit Ethernet interfaces 0/17 and 0/20:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet0/17 , gigabitethernet0/20
Switch(config-if-range)# no shutdown
Switch(config-if-range)#
*Oct 6 08:29:28: %LINK-3-UPDOWN: Interface GigabitEthernet0/17, changed
state to up
*Oct 6 08:29:28: %LINK-3-UPDOWN: Interface GigabitEthernet0/20, changed
state to up
```

If you enter multiple configuration commands while you are in interface-range mode, each command is executed as it is entered. The commands are not batched together and executed after you exit interface-range mode. If you exit interface-range configuration mode while the commands are being executed, some commands might not be executed on all interfaces in the range. Wait until the command prompt reappears before exiting interface-range configuration mode.

Configuring and Using Interface-Range Macros

You can create an interface-range macro to automatically select a range of interfaces for configuration. Before you can use the **macro** keyword in the **interface range macro** global configuration command string, you must use the **define interface-range** global configuration command to define the macro.

Beginning in privileged EXEC mode, follow these steps to define an interface-range macro:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	define interface-range <i>macro_name</i> <i>interface-range</i>	Define the interface-range macro, and save it in NVRAM. <ul style="list-style-type: none"> The <i>macro_name</i> is a 32-character maximum character string. A macro can contain up to five comma-separated interface ranges. You do not need to enter spaces before or after the comma. Each <i>interface-range</i> must consist of the same port type.
3.	interface range macro <i>macro_name</i>	Select the interface range to be configured by using the values saved in the interface-range macro called <i>macro_name</i> . You can now use the normal configuration commands to apply the configuration to all interfaces in the defined macro.
4.	end	Return to privileged EXEC mode.

Step	Command	Purpose
5.	show running-config include define	Show the defined interface-range macro configuration.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no define interface-range** *macro_name* global configuration command to delete a macro.

When using the **define interface-range** global configuration command, note these guidelines:

- Valid entries for *interface-range*:
 - vlan** *vlan-ID* - *vlan-ID*, where VLAN ID is from 1 to 4094
 - gigabitethernet** slot/{*first port*} - {*last port*}, where slot is **0**
 - port-channel** *port-channel-number* - *port-channel-number*, where *port-channel-number* is from 1 to 6.
- You must add a space between the interface numbers and the hyphen when entering an *interface-range*. For example, **gigabitethernet 0/1 - 5** is a valid range; **gigabitethernet 0/1-5** is not a valid range.
- The VLAN interfaces must have been configured with the **interface vlan** command. The **show running-config** privileged EXEC command output shows the configured VLAN interfaces. VLAN interfaces that do not appear by using the **show running-config** command cannot be used as *interface-ranges*.
- All interfaces in a range must be the same type; that is, all Gigabit Ethernet ports, all EtherChannel ports, or all VLANs, but you can combine multiple interface types in a macro.

This example shows how to define an interface-range macro named *enet_list* to select Gigabit Ethernet ports 17 to 20 and to verify the macro configuration:

```
Switch# configure terminal
Switch(config)# define interface-range enet_list gigabitethernet0/17 - 20
Switch(config)# end
Switch# show running-config | include define
define interface-range enet_list GigabitEthernet0/17 - 20
```

This example shows how to create a multiple-interface macro named *macro1*:

```
Switch# configure terminal
Switch(config)# define interface-range macro1 gi0/1 - 2, gi0/17 - 20
Switch(config)# end
Switch#
```

This example shows how to enter interface range configuration mode for the interface-range macro *enet_list*:

```
Switch# configure terminal
Switch(config)# interface range macro enet_list
Switch(config-if-range)#
```

This example shows how to delete the interface-range macro *enet_list* and to verify that it has been deleted.

```
Switch# configure terminal
```

```
Switch(config)# no define interface-range enet_list
Switch# show run | include define
```

Configuring Ethernet Interfaces

The switch supports these interface types:

- Physical ports—Switch ports, including access and trunk ports
- VLANs—Switch virtual interfaces (SVIs)
- Port-channels—EtherChannel of interfaces

These sections describe the default interface configuration and the optional features that you can configure on most physical interfaces:

- Default Ethernet Interface Configuration, on page 173
- Configuring Interface Speed and Duplex Mode, on page 174
- Configuring IEEE 802.3x Flow Control on Gigabit Ethernet Ports, on page 176
- Adding a Description for an Interface, on page 177

Default Ethernet Interface Configuration

Table 33 shows the Ethernet interface default configuration. For more details on the VLAN parameters listed in the table, see Chapter 13 “Configuring VLANs.” For details on controlling traffic to the port, see Chapter 16 “Configuring Port-Based Traffic Control.”

Table 33. Default Ethernet Interface Configuration

Feature	Default Setting
Operating mode	Layer 2
Allowed VLAN range	VLAN ID range is 2 to 4094 on the internal 1000 Mbps ports (ports 1 to 14) VLAN ID range is 1 to 1006 on the internal 100 Mbps management module ports (ports 15 and 16) VLAN ID range is 1 to 4094 on the external 10/100/1000 Mbps ports (ports 17 to 20)
Default VLAN	VLAN 2 in internal 1000 Mbps ports (ports 1 to 14) Native VLAN 1 on the internal 100 Mbps management module ports (ports 15 and 16) VLAN 1 on the external 10/100/1000 Mbps ports (ports 17 to 20)
Native VLAN (for 802.1Q trunks)	VLAN 2
VLAN trunking	Switchport mode dynamic desirable (supports DTP)
Port enable state	All ports are normally enabled. See the Installation Guide for information about changing this value.
Port description	blade n for the internal 1000 Mbps ports (ports 1 to 14) mgmt 1 or 2 for the internal 100 Mbps management module ports (ports 15 and 16) extern n for the external 10/100/1000 Mbps ports (ports 17 to 20)

Table 33. Default Ethernet Interface Configuration (continued)

Feature	Default Setting
Speed	1000 for the internal 1000 Mbps ports (ports 1 to 14) 100 for the internal 100 Mbps management module ports (ports 15 and 16) Autonegotiate for the external 10/100/1000 Mbps ports (ports 17 to 20) The speed on the internal ports is non-configurable.
Duplex mode	Full duplex for the internal 1000 Mbps ports (ports 1 to 14) Full duplex for the internal 100 Mbps management module ports (ports 15 and 16) Autonegotiate for the external 10/100/1000 Mbps ports (ports 17 to 20) The duplex mode on internal ports is non-configurable.
Flow control	Flow control is set to <i>off</i> for receive and <i>desired</i> for send for Gigabit Ethernet ports.
EtherChannel (PAgP) and Link Aggregation Control Protocol (LACP)	Disabled on all Ethernet ports. See Chapter 25 “Configuring EtherChannels.”
Broadcast, multicast, and unicast storm control	Disabled. See the “Default Storm Control Configuration” section on page 316.
Protected port	Disabled. See the “Configuring Protected Ports” section on page 317.
Port security	Disabled. See the “Default Port Security Configuration” section on page 320.
Port Fast	Enabled

Configuring Interface Speed and Duplex Mode

The external 10/100/1000 Mbps ports (ports 17 to 20) autonegotiate in speed and duplex mode. You can change the speed and duplex settings of the external ports. The internal Gigabit Ethernet ports (ports 1 to 14) operate at 1000 Mbps, full duplex only. The internal 100 Mbps management module ports (ports 15 and 16) operate at 100 Mbps, full duplex only. You cannot configure the speed and duplex mode on the internal ports.

In full-duplex mode, two stations can send and receive at the same time. When packets can flow in both directions simultaneously, effective Ethernet bandwidth doubles to 2 Gbps for Gigabit Ethernet interfaces. Full-duplex communication is often an effective solution to collisions, which are major constrictions in Ethernet networks. Normally, 10-Mbps ports operate in half-duplex mode, which means that stations can either receive or send.

These sections describe how to configure the interface speed and duplex mode:

- Configuration Guidelines, on page 175
- Setting the Interface Speed and Duplex Parameters on a Switch Port, on page 175

Configuration Guidelines

When configuring an interface speed and duplex mode, note these guidelines:

- Ethernet ports set to 1000 Mbps are always be set to full duplex.
- Gigabit Ethernet ports that do not match the settings of an attached device can lose connectivity and do not generate statistics.
- If both ends of the line support autonegotiation, we highly recommend the default setting of **autonegotiation**.
- When connecting an interface to a 100BASE-T device that does not autonegotiate, set the speed to a non-auto value (for example, **nonegotiate**) and set the duplex mode to full or half to match the device. The speed value and duplex mode must be explicitly set.
- When connecting an interface to a Gigabit Ethernet device that does not autonegotiate, disable autonegotiation on the switch and set the duplex and flow control parameters to be compatible with the remote device.

Caution: Changing the interface speed and duplex mode configuration might shut down and re-enable the interface during the reconfiguration.

Setting the Interface Speed and Duplex Parameters on a Switch Port

Beginning in privileged EXEC mode, follow these steps to set the speed and duplex mode for a physical interface on a non-LRE switch:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode and the physical interface identification.
3.	speed {10 100 1000 auto nonegotiate}	Enter the appropriate speed parameter for the interface, or enter auto or nonegotiate . Note: This option is available only on the external 10/100/1000 Mbps ports. It is not available on the internal switch ports.
4.	duplex {auto full half}	Enter the duplex parameter for the interface. Note: This option is available only on the external 10/100/1000 Mbps ports. It is not available on the internal switch ports.
5.	end	Return to privileged EXEC mode.
6.	show interfaces <i>interface-id</i>	Display the interface speed and duplex mode configuration.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no speed** and **no duplex** interface configuration commands to return the interface to the default speed and duplex settings (autonegotiate). To return all interface settings to the defaults, use the **default interface** *interface-id* interface configuration command.

This example shows how to set the interface speed to 10 Mbps and the duplex mode to half on Gigabit Ethernet interface 0/17 and to verify the configuration:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# speed 10
```

```

Switch(config-if)# duplex half
Switch(config)# end
Switch# show running-config
Building configuration...

Current configuration : 1954 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname Switch
!
<output truncated>
!
interface gigabitethernet0/17
  switchport mode trunk
  no ip address
  duplex half
  speed 10
!
<output truncated>

```

Configuring IEEE 802.3x Flow Control on Gigabit Ethernet Ports

Flow control is supported only on the 10/100/1000 Mbps ports. Flow control enables connected Gigabit Ethernet ports to control traffic rates during congestion by allowing congested nodes to pause link operation at the other end. If one port experiences congestion and cannot receive any more traffic, it notifies the other port to stop sending until the condition clears. When the local device detects any congestion at its end, it can notify the link partner or the remote device of the congestion by sending a pause frame. Upon receipt of a pause frame, the remote device stops sending any data packets, which prevents any loss of data packets during the congestion period.

Note: We strongly recommend that you do not configure IEEE 802.3x flow control when quality of service (QoS) is configured on the switch. Before configuring flow control on an interface, make sure to return to the default QoS settings listed in the “Default Standard QoS Configuration” section on page 422.

Flow control can be implemented in two forms, symmetric and asymmetric. The symmetric implementation is suitable for point-to-point links, and asymmetric is suitable for hub-to-end node connections, where it is desirable for the hub to pause the end system, but not vice-versa. You use the **flowcontrol** interface configuration command to set the interface’s ability to **receive** and **send** pause frames to **on**, **off**, or **desired**. The default state for Gigabit Ethernet ports is **receive off** and **send desired**. The default state for Fast Ethernet ports is **receive off** and **send off**.

These rules apply to flow control settings on the device:

- **receive on** (or **desired**) and **send on**: Flow control operates in both directions; both the local and the remote devices can send pause frames to show link congestion.
- **receive on** (or **desired**) and **send desired**: The port can receive pause frames and can send pause frames if the attached device supports flow control.

- **receive on** (or **desired**) and **send off**: The port cannot send pause frames but can operate with an attached device that is required to or can send pause frames; the port can receive pause frames.
- **receive off** and **send on**: The port sends pause frames if the remote device supports flow control but cannot receive pause frames from the remote device.
- **receive off** and **send desired**: The port cannot receive pause frames but can send pause frames if the attached device supports flow control.
- **receive off** and **send off**: Flow control does not operate in either direction. In case of congestion, no indication is given to the link partner, and no pause frames are sent or received by either device.

Note: For details about the command settings and the resulting flow control resolution on local and remote ports, refer to the **flowcontrol** interface configuration command in the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference* for this release.

Beginning in privileged EXEC mode, follow these steps to configure flow control on an interface:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode
2.	interface <i>interface-id</i>	Enter interface configuration mode and the physical interface to be configured.
3.	flowcontrol { receive send } { on off desired }	Configure the flow control mode for the port.
4.	end	Return to privileged EXEC mode.
5.	show interfaces <i>interface-id</i>	Verify the interface flow control settings.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable flow control, use the **flowcontrol receive off** and **flowcontrol send off** interface configuration commands.

This example shows how to turn off all flow control on Gigabit Ethernet interface 0/17 and to display the results:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# flowcontrol receive off
Switch(config-if)# flowcontrol send off
Switch(config-if)# end
Switch# show running-config
```

Adding a Description for an Interface

You can add a description about an interface to help you remember its function. The description appears in the output of these commands: **show configuration**, **show running-config**, and **show interfaces**.

Beginning in privileged EXEC mode, follow these steps to add a description for an interface:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode
2.	interface <i>interface-id</i>	Enter interface configuration mode, and enter the interface for which you are adding a description.
3.	description <i>string</i>	Add a description (up to 240 characters) for an interface.
4.	end	Return to privileged EXEC mode.
5.	show interfaces <i>interface-id</i> description or show running-config	Verify your entry.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no description** interface configuration command to delete the description.

This example shows how to add a description on Gigabit Ethernet interface 0/17 and to verify the description:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# description Connects to Marketing
Switch(config-if)# end
Switch# show interfaces gigabitethernet0/17 description
Interface Status      Protocol Description
Gi0/17    up                down     Connects to Marketing
```

Monitoring and Maintaining the Interfaces

You can perform the tasks in these sections to monitor and maintain interfaces:

- Monitoring Interface and Controller Status, on page 178
- Clearing and Resetting Interfaces and Counters, on page 180
- Shutting Down and Restarting the Interface, on page 181

Monitoring Interface and Controller Status

Commands entered at the privileged EXEC prompt display information about the interface, including the version of the software and the hardware, the controller status, and statistics about the interfaces. Table 34 lists some of these interface monitoring commands. (You can display the full list of **show** commands by using the **show ?** command at the privileged EXEC prompt.) These commands are fully described in the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference*.

Table 34. *show* Commands for Interfaces

Command	Purpose
show interfaces [<i>interface-id</i>]	Display the status and configuration of all interfaces or a specific interface.
show interfaces [<i>interface-id</i>] capabilities [module { <i>module-number</i> }]	Display the capabilities of an interface. If you do not specify a module, the capabilities for all ports on the switch are displayed.
show interfaces <i>interface-id</i> status [err-disabled]	Display interface status or a list of interfaces in error-disabled state.
show interfaces [<i>interface-id</i>] switchport	Display administrative and operational status of switching (nonrouting) ports.
show interfaces [<i>interface-id</i>] description	Display the description configured on an interface or all interfaces and the interface status.
show ip interface [<i>interface-id</i>]	Display the usability status of all interfaces configured for IP or the specified interface.
show running-config interface [<i>interface-id</i>]	Display the running configuration in RAM for the interface.
show version	Display the hardware configuration, software version, the names and sources of configuration files, and the boot images.

This example shows how to display the status of switching ports:

```
Switch# show interfaces switchport
Name: Gi0/1
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: down
Administrative Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 2 (operational)
Trunking Native Mode VLAN: 2 (operational)
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk private VLANs: none
Operational private-vlan: none
Trunking VLANs Enabled: 2-4094
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL
Protected: false

Appliance trust: none
.
.
.
<output truncated>
```

This example shows how to display the running configuration of Gigabit Ethernet interface 17:

```
Switch# show running-config interface gigabitethernet0/17
Building configuration...

Current configuration : 156 bytes
!
interface GigabitEthernet0/17
  description external
  switchport access vlan 2
  switchport trunk native vlan 2
  ip access-group SecWiz_Gi0_1_out_ip in
end
```

For additional examples of the **show interfaces** privileged EXEC command output, refer to the command reference for this release.

Clearing and Resetting Interfaces and Counters

Table 35 lists the privileged EXEC mode **clear** commands that you can use to clear counters and reset interfaces.

Table 35. Clear Commands for Interfaces

Command	Purpose
clear counters [<i>interface-id</i>]	Clear interface counters.
clear interface <i>interface-id</i>	Reset the hardware logic on an interface.
clear line [<i>number</i> console 0 vty number]	Reset the hardware logic on an asynchronous serial line.

To clear the interface counters shown by the **show interfaces** privileged EXEC command, use the **clear counters** privileged EXEC command. The **clear counters** command clears all current interface counters from the interface unless optional arguments are specified to clear only a specific interface type from a specific interface number.

Note: The **clear counters** privileged EXEC command does not clear counters retrieved by using Simple Network Management Protocol (SNMP), but only those seen with the **show interfaces** privileged EXEC command output.

This example shows how to clear and reset the counters on Gigabit Ethernet interface 0/17:

```
Switch# clear counters gigabitethernet0/17
Clear "show interface" counters on this interface [confirm] y
Switch#
*Sep 30 08:42:55: %CLEAR-5-COUNTERS: Clear counter on interface
gigabitethernet0/17
by vty1 (171.69.115.10)
```

Use the **clear interface** or **clear line** privileged EXEC command to clear and reset an interface or serial line. Under most circumstances, you do not need to clear the hardware logic on interfaces or serial lines.

This example shows how to clear and reset Gigabit Ethernet interface 0/17:

```
Switch# clear interface gigabitethernet0/17
```

Shutting Down and Restarting the Interface

Shutting down an interface disables all functions on the specified interface and marks the interface as unavailable on all monitoring command displays. This information is communicated to other network servers through all dynamic routing protocols. The interface is not mentioned in any routing updates.

Beginning in privileged EXEC mode, follow these steps to shut down an interface:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface {vlan <i>vlan-id</i> } {{ gigabitethernet gigabitethernet } <i>interface-id</i> } { port-channel <i>port-channel-number</i> }	Select the interface to be configured.
3.	shutdown	Shut down an interface.
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entry.

Use the **no shutdown** interface configuration command to restart the interface. This example shows how to shut down Gigabit Ethernet interface 0/17:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# shutdown
Switch(config-if)#
*Sep 30 08:33:47: %LINK-5-CHANGED: Interface GigabitEthernet0/17, changed
state to a administratively down
```

This example shows how to re-enable Gigabit Ethernet interface 0/17:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# no shutdown
Switch(config-if)#
*Sep 30 08:36:00: %LINK-3-UPDOWN: Interface GigabitEthernet0/17, changed
state to up
```

To verify that an interface is disabled, enter the **show interfaces** privileged EXEC command. A disabled interface is shown as *administratively down* in the **show interfaces** command output.

Chapter 10. Configuring STP

This chapter describes how to configure the Spanning Tree Protocol (STP) on port-based VLANs on your switch. The switch uses the per-VLAN spanning-tree plus (PVST+) protocol based on the IEEE 802.1D standard and Cisco proprietary extensions, or it can use the rapid per-VLAN spanning-tree plus (rapid-PVST+) protocol based on the IEEE 802.1w standard.

For information about the Multiple Spanning Tree Protocol (MSTP) and how to map multiple VLANs to the same spanning-tree instance, see Chapter 11 “Configuring MSTP.”

For information about other spanning-tree features such as Port Fast, UplinkFast, root guard, and so forth, see Chapter 12 “Configuring Optional Spanning-Tree Features.”

Note: For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

This chapter consists of these sections:

- Understanding Spanning-Tree Features, on page 183
- Configuring Spanning-Tree Features, on page 192
- Displaying the Spanning-Tree Status, on page 203

Understanding Spanning-Tree Features

These sections describe how basic spanning-tree features work:

- STP Overview, on page 183
- Spanning-Tree Topology and BPDUs, on page 184
- Bridge ID, Switch Priority, and Extended System ID, on page 185
- Spanning-Tree Interface States, on page 186
- How a Switch or Port Becomes the Root Switch or Root Port, on page 188
- Spanning Tree and Redundant Connectivity, on page 189
- Spanning-Tree Address Management, on page 190
- Accelerated Aging to Retain Connectivity, on page 190
- Spanning-Tree Modes and Protocols, on page 190
- Supported Spanning-Tree Instances, on page 191
- Spanning-Tree Interoperability and Backward Compatibility, on page 191
- STP and IEEE 802.1Q Trunks, on page 191
- Spanning Tree Considerations for Cisco Systems Intelligent Gigabit Ethernet Switch Modules, on page 192

For configuration information, see the “Configuring Spanning-Tree Features” section on page 192.

For information about optional spanning-tree features, see Chapter 12 “Configuring Optional Spanning-Tree Features.”

STP Overview

STP is a Layer 2 link management protocol that provides path redundancy while preventing loops in the network. For a Layer 2 Ethernet network to function properly,

only one active path can exist between any two stations. Multiple active paths among end stations cause loops in the network. If a loop exists in the network, end stations might receive duplicate messages. Switches might also learn end-station MAC addresses on multiple Layer 2 interfaces. These conditions result in an unstable network. Spanning-tree operation is transparent to end stations, which cannot detect whether they are connected to a single LAN segment or a switched LAN of multiple segments.

The STP uses a spanning-tree algorithm to select one switch of a redundantly connected network as the root of the spanning tree. The algorithm calculates the best loop-free path through a switched Layer 2 network by assigning a role to each port based on the role of the port in the active topology:

- Root—A forwarding port elected for the spanning-tree topology
- Designated—A forwarding port elected for every switched LAN segment
- Alternate—A blocked port providing an alternate path to the root port in the spanning tree
- Backup—A blocked port in a loopback configuration

Switches that have ports with these assigned roles are called root or designated switches.

Spanning tree forces redundant data paths into a standby (blocked) state. If a network segment in the spanning tree fails and a redundant path exists, the spanning-tree algorithm recalculates the spanning-tree topology and activates the standby path. Switches send and receive spanning-tree frames, called bridge protocol data units (BPDUs), at regular intervals. The switches do not forward these frames but use them to construct a loop-free path. BPDUs contain information about the sending switch and its ports, including switch and MAC addresses, switch priority, port priority, and path cost. Spanning tree uses this information to elect the root switch and root port for the switched network and the root port and designated port for each switched segment.

When two interfaces on a switch are part of a loop, the spanning-tree port priority and path cost settings determine which interface is put in the forwarding state and which is put in the blocking state. The spanning-tree port priority value represents the location of an interface in the network topology and how well it is located to pass traffic. The path cost value represents the media speed.

Spanning-Tree Topology and BPDUs

The stable, active spanning-tree topology of a switched network is determined by these elements:

- The unique bridge ID (switch priority and MAC address) associated with each VLAN on each switch
- The spanning-tree path cost to the root switch
- The port identifier (port priority and MAC address) associated with each Layer 2 interface

When the switches in a network are powered up, each functions as the root switch. Each switch sends a configuration BPDU through all of its ports. The BPDUs communicate and compute the spanning-tree topology. Each configuration BPDU contains this information:

- The unique bridge ID of the switch that the sending switch identifies as the root switch
- The spanning-tree path cost to the root

- The bridge ID of the sending switch
- Message age
- The identifier of the sending interface
- Values for the hello, forward-delay, and max-age protocol timers

When a switch receives a configuration BPDU that contains *superior* information (lower bridge ID, lower path cost, and so forth), it stores the information for that port. If this BPDU is received on the root port of the switch, the switch also forwards it with an updated message to all attached LANs for which it is the designated switch.

If a switch receives a configuration BPDU that contains *inferior* information to that currently stored for that port, it discards the BPDU. If the switch is a designated switch for the LAN from which the inferior BPDU was received, it sends that LAN a BPDU containing the up-to-date information stored for that port. In this way, inferior information is discarded, and superior information is propagated on the network.

A BPDU exchange results in these actions:

- One switch in the network is elected as the root switch (the logical center of the spanning-tree topology in a switched network).
For each VLAN, the switch with the highest switch priority (the lowest numerical priority value) is elected as the root switch. If all switches are configured with the default priority (32768), the switch with the lowest MAC address in the VLAN becomes the root switch. The switch priority value occupies the most significant bits of the bridge ID, as shown in Table 36 on page 186.
- A root port is selected for each switch (except the root switch). This port provides the best path (lowest cost) when the switch forwards packets to the root switch.
- The shortest distance to the root switch is calculated for each switch based on the path cost.
- A designated switch for each LAN segment is selected. The designated switch incurs the lowest path cost when forwarding packets from that LAN to the root switch. The port through which the designated switch is attached to the LAN is called the designated port.
- Interfaces included in the spanning-tree instance are selected. Root ports and designated ports are put in the forwarding state.
- All paths that are not needed to reach the root switch from anywhere in the switched network are placed in the spanning-tree blocking mode.

Bridge ID, Switch Priority, and Extended System ID

The IEEE 802.1D standard requires that each switch has a unique bridge identifier (bridge ID), which determines the selection of the root switch. Because each VLAN is considered as a different *logical bridge* with PVST+ and rapid PVST+, the same switch must have as many different bridge IDs as VLANs configured on it. Each VLAN on the switch has a unique 8-byte bridge ID; the two most-significant bytes are used for the switch priority, and the remaining six bytes are derived from the switch MAC address.

The switches support the 802.1t spanning-tree extensions. Some of the bits previously used for the switch priority are now used as the VLAN identifier. The result is that fewer MAC addresses are reserved for the switch, and a larger range of VLAN IDs can be supported, all while maintaining the uniqueness of the bridge ID. As shown in Table 36, the two bytes previously used for the switch priority are reallocated into a 4-bit priority value and a 12-bit extended system ID value equal to the VLAN ID. In earlier releases, the switch priority is a 16-bit value.

Table 36. Switch Priority Value and Extended System ID

Switch Priority Value				Extended System ID (Set Equal to the VLAN ID)											
Bit 16	Bit 15	Bit 14	Bit 13	Bit 12	Bit 11	Bit 10	Bit 9	Bit 8	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1
32768	16384	8192	4096	2048	1024	512	256	128	64	32	16	8	4	2	1

Spanning tree uses the extended system ID, the switch priority, and the allocated spanning-tree MAC address to make the bridge ID unique for each VLAN. With earlier releases, spanning tree used one MAC address per VLAN to make the bridge ID unique for each VLAN.

Support for the extended system ID affects how you manually configure the root switch, the secondary root switch, and the switch priority of a VLAN. For example, when you change the switch priority value, you change the probability that the switch will be elected as the root switch. Configuring a higher value decreases the probability; a lower value increases the probability. For more information, see the “Configuring the Root Switch” section on page 196, the “Configuring a Secondary Root Switch” section on page 197, and the “Configuring the Switch Priority of a VLAN” section on page 201.

Spanning-Tree Interface States

Propagation delays can occur when protocol information passes through a switched LAN. As a result, topology changes can take place at different times and at different places in a switched network. When an interface transitions directly from nonparticipation in the spanning-tree topology to the forwarding state, it can create temporary data loops. Interfaces must wait for new topology information to propagate through the switched LAN before starting to forward frames. They must allow the frame lifetime to expire for forwarded frames that have used the old topology.

Each Layer 2 interface on a switch using spanning tree exists in one of these states:

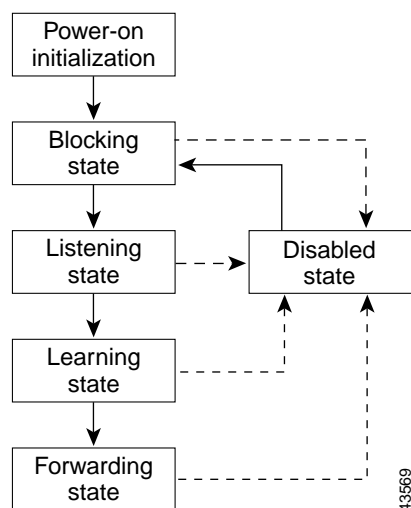
- Blocking—The interface does not participate in frame forwarding.
- Listening—The first transitional state after the blocking state when the spanning tree determines that the interface should participate in frame forwarding.
- Learning—The interface prepares to participate in frame forwarding.
- Forwarding—The interface forwards frames.
- Disabled—The interface is not participating in spanning tree because of a shutdown port, no link on the port, or no spanning-tree instance running on the port.

An interface moves through these states:

- From initialization to blocking
- From blocking to listening or to disabled
- From listening to learning or to disabled
- From learning to forwarding or to disabled
- From forwarding to disabled

Figure 35 illustrates how an interface moves through the states.

Figure 35. Spanning-Tree Interface States



When you power up the switch, spanning tree is enabled by default, and every interface in the switch, VLAN, or network goes through the blocking state and the transitory states of listening and learning. Spanning tree stabilizes each interface at the forwarding or blocking state.

When the spanning-tree algorithm places a Layer 2 interface in the forwarding state, this process occurs:

1. The interface is in the listening state while spanning tree waits for protocol information to transition the interface to the blocking state.
2. While spanning tree waits the forward-delay timer to expire, it moves the interface to the learning state and resets the forward-delay timer.
3. In the learning state, the interface continues to block frame forwarding as the switch learns end-station location information for the forwarding database.
4. When the forward-delay timer expires, spanning tree moves the interface to the forwarding state, where both learning and frame forwarding are enabled.

Blocking State

A Layer 2 interface in the blocking state does not participate in frame forwarding. After initialization, a BPDU is sent to each interface in the switch. A switch initially functions as the root until it exchanges BPDUs with other switches. This exchange establishes which switch in the network is the root or root switch. If there is only one switch in the network, no exchange occurs, the forward-delay timer expires, and the interfaces move to the listening state. An interface always enters the blocking state after switch initialization.

An interface in the blocking state performs as follows:

- Discards frames received on the port
- Discards frames switched from another interface for forwarding
- Does not learn addresses
- Receives BPDUs

Listening State

The listening state is the first state a Layer 2 interface enters after the blocking state. The interface enters this state when the spanning tree determines that the interface should participate in frame forwarding.

An interface in the listening state performs as follows:

- Discards frames received on the port
- Discards frames switched from another interface for forwarding
- Does not learn addresses
- Receives BPDUs

Learning State

A Layer 2 interface in the learning state prepares to participate in frame forwarding. The interface enters the learning state from the listening state.

An interface in the learning state performs as follows:

- Discards frames received on the port
- Discards frames switched from another interface for forwarding
- Learns addresses
- Receives BPDUs

Forwarding State

A Layer 2 interface in the forwarding state forwards frames. The interface enters the forwarding state from the learning state.

An interface in the forwarding state performs as follows:

- Receives and forwards frames received on the port
- Forwards frames switched from another port
- Learns addresses
- Receives BPDUs

Disabled State

A Layer 2 interface in the disabled state does not participate in frame forwarding or in the spanning tree. An interface in the disabled state is nonoperational.

A disabled interface performs as follows:


- Discards frames received on the port
- Discards frames switched from another interface for forwarding
- Does not learn addresses
- Does not receive BPDUs

How a Switch or Port Becomes the Root Switch or Root Port

If all switches in a network are enabled with default spanning-tree settings, the switch with the lowest MAC address becomes the root switch. In Figure 36, Switch A is elected as the root switch because the switch priority of all the switches is set to the default (32768) and Switch A has the lowest MAC address. However, because of traffic patterns, number of forwarding interfaces, or link types, Switch A might not be

the ideal root switch. By increasing the priority (lowering the numerical value) of the ideal switch so that it becomes the root switch, you force a spanning-tree recalculation to form a new topology with the ideal switch as the root.

Figure 36. Spanning-Tree Topology



RP = Root Port
DP = Designated Port


When the spanning-tree topology is calculated based on default parameters, the path between source and destination end stations in a switched network might not be ideal. For instance, connecting higher-speed links to an interface that has a higher number than the root port can cause a root-port change. The goal is to make the fastest link the root port.

For example, assume that one port on Switch B is a Gigabit Ethernet link and that another port on Switch B (a 100 Mbps link) is the root port. Network traffic might be more efficient over the Gigabit Ethernet link. By changing the spanning-tree port priority on the Gigabit Ethernet interface to a higher priority (lower numerical value) than the root port, the Gigabit Ethernet interface becomes the new root port.

Spanning Tree and Redundant Connectivity

You can create a redundant backbone with spanning tree by connecting two switch interfaces to another device or to two different devices. Spanning tree automatically disables one interface but enables it if the other one fails, as shown in Figure 37. If one link is high-speed and the other is low-speed, the low-speed link is always disabled. If the speeds are the same, the port priority and port ID are added together, and spanning tree disables the link with the lowest value.

Figure 37. Spanning Tree and Redundant Connectivity



You can also create redundant links between switches by using EtherChannel groups. For more information, see Chapter 25 “Configuring EtherChannels.”

Spanning-Tree Address Management

IEEE 802.1D specifies 17 multicast addresses, ranging from 0x00180C2000000 to 0x0180C2000010, to be used by different bridge protocols. These addresses are static addresses that cannot be removed.

Regardless of the spanning-tree state, the switch receives but does not forward packets destined for addresses between 0x0180C2000000 and 0x0180C200000F.

If spanning tree is enabled, the switch CPU receives packets destined for 0x0180C2000000 and 0x0180C2000010. If spanning-tree is disabled, the switch forwards those packets as unknown multicast addresses.

Accelerated Aging to Retain Connectivity

The default for aging dynamic addresses is 5 minutes, the default setting of the **mac-address-table aging-time** global configuration command. However, a spanning-tree reconfiguration can cause many station locations to change. Because these stations could be unreachable for 5 minutes or more during a reconfiguration, the address-aging time is accelerated so that station addresses can be dropped from the address table and then relearned. The accelerated aging is the same as the forward-delay parameter value (**spanning-tree vlan *vlan-id* forward-time *seconds*** global configuration command) when the spanning tree reconfigures.

Because each VLAN is a separate spanning-tree instance, the switch accelerates aging on a per-VLAN basis. A spanning-tree reconfiguration on one VLAN can cause the dynamic addresses learned on that VLAN to be subject to accelerated aging. Dynamic addresses on other VLANs can be unaffected and remain subject to the aging interval entered for the switch.

Spanning-Tree Modes and Protocols

The switch supports these spanning-tree modes and protocols:

- **PVST+**—This spanning-tree mode is based on the IEEE 802.1D standard and Cisco proprietary extensions. The PVST+ runs on each VLAN on the switch up to the maximum supported, ensuring that each has a loop-free path through the network.

The PVST+ provides Layer 2 load balancing for the VLAN on which it runs. You can create different logical topologies by using the VLANs on your network to ensure that all of your links are used but that no one link is oversubscribed. Each instance of PVST+ on a VLAN has a single root switch. This root switch propagates the spanning-tree information associated with that VLAN to all other switches in the network. Because each switch has the same information about the network, this process ensures that the network topology is maintained.

- **Rapid PVST+**—This spanning-tree mode is the same as PVST+ except that it uses a rapid convergence based on the IEEE 802.1w standard. It is the default spanning-tree mode used on all Ethernet, Fast Ethernet, and Gigabit Ethernet port-based VLANs. To provide rapid convergence, the rapid PVST+ immediately deletes dynamically learned MAC address entries on a per-port basis upon receiving a topology change. By contrast, PVST+ uses a short aging time for dynamically learned MAC address entries.

The rapid PVST+ uses the same configuration as PVST+ (except where noted), and the switch needs only minimal extra configuration. The benefit of rapid PVST+

is that you can migrate a large PVST+ install base to rapid PVST+ without having to learn the complexities of the MSTP configuration and without having to re-provision your network. In rapid-PVST+ mode, each VLAN runs its own spanning-tree instance up to the maximum supported.

- **MSTP**—This spanning-tree mode is based on the IEEE 802.1s standard. You can map multiple VLANs to the same spanning-tree instance, which reduces the number of spanning-tree instances required to support a large number of VLANs. The MSTP runs on top of the RSTP (based on IEEE 802.1w), which provides for rapid convergence of the spanning tree by eliminating the forward delay and by quickly transitioning root ports and designated ports to the forwarding state. You cannot run MSTP without RSTP.

The most common initial deployment of MSTP is in the backbone and distribution layers of a Layer 2 switched network. For more information, see Chapter 11 “Configuring MSTP.” For information about the number of supported spanning-tree instances, see the next section.

Supported Spanning-Tree Instances

In PVST+ or rapid-PVST+ mode, the switch supports up to 64 spanning-tree instances.

In MSTP mode, the switch supports up to 16 MST instances. The number of VLANs that can be mapped to a particular MST instance is unlimited.

For information about how spanning tree interoperates with the VLAN Trunking Protocol (VTP), see the “Spanning-Tree Configuration Guidelines” section on page 193.

Spanning-Tree Interoperability and Backward Compatibility

Table 37 lists the interoperability and compatibility among the supported spanning-tree modes in a network.

Table 37. PVST+, MSTP, and Rapid-PVST+ Interoperability

	PVST+	MSTP	Rapid PVST+
PVST+	Yes	Yes (with restrictions)	Yes (reverts to PVST+)
MSTP	Yes (with restrictions)	Yes	Yes (reverts to PVST+)
Rapid PVST+	Yes (reverts to PVST+)	Yes (reverts to PVST+)	Yes

In a mixed MSTP and PVST+ network, the common spanning-tree (CST) root must be inside the MST backbone, and a PVST+ switch cannot connect to multiple MST regions.

When a network contains switches running rapid PVST+ and switches running PVST+, we recommend that the rapid-PVST+ switches and PVST+ switches be configured for different spanning-tree instances. In the rapid-PVST+ spanning-tree instances, the root switch must be a rapid-PVST+ switch. In the PVST+ instances, the root switch must be a PVST+ switch. The PVST+ switches should be at the edge of the network.

STP and IEEE 802.1Q Trunks

The IEEE 802.1Q standard for VLAN trunks imposes some limitations on the spanning-tree strategy for a network. The standard requires only one spanning-tree instance for *all* VLANs allowed on the trunks. However, in a network of Cisco switches

connected through 802.1Q trunks, the switches maintain one spanning-tree instance for *each* VLAN allowed on the trunks.

When you connect a Cisco switch to a non-Cisco device through an 802.1Q trunk, the Cisco switch uses PVST+ to provide spanning-tree interoperability. If rapid PVST+ is enabled, the switch uses it instead of PVST+. The switch combines the spanning-tree instance of the 802.1Q VLAN of the trunk with the spanning-tree instance of the non-Cisco 802.1Q switch.

However, all PVST+ or rapid-PVST+ information is maintained by Cisco switches separated by a cloud of non-Cisco 802.1Q switches. The non-Cisco 802.1Q cloud separating the Cisco switches is treated as a single trunk link between the switches.

The external spanning-tree behavior on access ports and trunk ports is not affected by PVST+ or rapid PVST+.

For more information on 802.1Q trunks, see Chapter 13 “Configuring VLANs.”

Spanning Tree Considerations for Cisco Systems Intelligent Gigabit Ethernet Switch Modules

A port-blocking filter exists between the switch external ports and the switch internal management module ports. This filter prevents operational traffic (such as unicast, multicast, and broadcast traffic) entering a switch external port from being forwarded to the management module, and from the management module to the external ports.

However, STP does not recognize that this filter exists. During topology discovery, STP incorrectly perceives that an external port is forwarding operational traffic to the internal management module ports and that a Layer 2 loop exists. STP changes the state of the internal management module port to blocked state. This action is acceptable for operational traffic, but not for management (non-operational) traffic.

The default path cost value on the switch internal management module ports is 100. The intent is to block operational traffic from being forwarded to the management module through any external port in the non-management VLAN. STP will see the cost of the management module port as the most expensive and block it. We do not recommend using the management module ports to carry operational traffic. This does not apply to the management VLAN on the management module ports. The switch prevents STP from blocking the management VLAN on the management module ports. STP blocking of VLANs on the management module ports is permitted for non-management VLANs only.

Configuring Spanning-Tree Features

These sections describe how to configure spanning-tree features:

- Default Spanning-Tree Configuration, on page 193
- Spanning-Tree Configuration Guidelines, on page 193
- Changing the Spanning-Tree Mode, on page 194 (required)
- Disabling Spanning Tree, on page 195 (optional)
- Configuring the Root Switch, on page 196 (optional)
- Configuring a Secondary Root Switch, on page 197 (optional)
- Configuring the Port Priority, on page 198 (optional)
- Configuring the Path Cost, on page 199 (optional)
- Configuring the Switch Priority of a VLAN, on page 201 (optional)

- Configuring Spanning-Tree Timers, on page 201 (optional)

Default Spanning-Tree Configuration

Table 38 shows the default spanning-tree configuration.

Table 38. Default Spanning-Tree Configuration

Feature	Default Setting
Enable state	<p>Enabled on VLAN 1 (default management VLAN for the management module).</p> <p>Enabled on VLAN 2 (default operational traffic VLAN for the external 10/100/1000 Mbps ports and the internal Gigabit Ethernet ports).</p> <p>For more information, see the “Supported Spanning-Tree Instances” section on page 191.</p>
Spanning-tree mode	Rapid PVST+. (PVST+ and MSTP are disabled.)
Switch priority	32768.
Spanning-tree port priority (configurable on a per-interface basis)	128.
Spanning-tree port cost (configurable on a per-interface basis)	<p>1000 Mbps: 4.</p> <p>100 Mbps: 19.</p> <p>10 Mbps: 100.</p>
Spanning-tree VLAN port priority (configurable on a per-VLAN basis)	128.
Spanning-tree VLAN port cost (configurable on a per-VLAN basis)	<p>1000 Mbps: 4.</p> <p>100 Mbps: 19.</p> <p>10 Mbps: 100.</p>
Spanning-tree timers	<p>Hello time: 2 seconds.</p> <p>Forward-delay time: 15 seconds.</p> <p>Maximum-aging time: 20 seconds.</p>

Spanning-Tree Configuration Guidelines

If more VLANs are defined in the VTP than there are spanning-tree instances, you can enable PVST+ or rapid PVST+ on only 64 VLANs. If the number of VLANs exceeds 64, we recommend that you enable the MSTP to map multiple VLANs to a single spanning-tree instance. For more information, see the Chapter 11 “Configuring MSTP.”

If 64 instances of spanning tree are already in use, you can disable spanning tree on one of the VLANs and then enable it on the VLAN where you want it to run. Use the **no spanning-tree vlan** *vlan-id* global configuration command to disable spanning tree on a specific VLAN, and use the **spanning-tree vlan** *vlan-id* global configuration command to enable spanning tree on the desired VLAN.

Caution: Switches that are not running spanning tree still forward BPDUs that they receive so that the other switches on the VLAN that have a running spanning-tree instance can break loops. Therefore, spanning tree must be running on enough switches to break all the loops in the network; for example, at least one switch on each loop in the VLAN must be running spanning tree. It is not absolutely necessary to run spanning tree on all switches in the VLAN; however, if you are running spanning tree only on a minimal set of switches, an incautious change to the network that introduces another loop into the VLAN can result in a broadcast storm.

Note: If you have already used all available spanning-tree instances on your switch, adding another VLAN anywhere in the VTP domain creates a VLAN that is not running spanning tree on that switch. If you have the default allowed list on the trunk ports of that switch, the new VLAN is carried on all trunk ports. Depending on the topology of the network, this could create a loop in the new VLAN that will not be broken, particularly if there are several adjacent switches that have all run out of spanning-tree instances. You can prevent this possibility by setting up allowed lists on the trunk ports of switches that have used up their allocation of spanning-tree instances. Setting up allowed lists is not necessary in many cases and can make it more labor-intensive to add another VLAN to the network.

Spanning-tree commands determine the configuration of VLAN spanning-tree instances. You create a spanning-tree instance when you assign an interface to a VLAN. The spanning-tree instance is removed when the last interface is moved to another VLAN. You can configure switch and port parameters before a spanning-tree instance is created; these parameters are applied when the spanning-tree instance is created.

The switch supports PVST+, rapid PVST+, and MSTP, but only one version can be active at any time. (For example, all VLANs run PVST+, all VLANs run rapid PVST+, or all VLANs run MSTP.) For information about the different spanning-tree modes and how they interoperate, see the “Spanning-Tree Interoperability and Backward Compatibility” section on page 191.

Changing the Spanning-Tree Mode

The switch supports three spanning-tree modes: PVST+, rapid PVST+, or MSTP. By default, the switch runs the PVST+ protocol.

Beginning in privileged EXEC mode, follow these steps to change the spanning-tree mode. If you want to enable a mode that is different from the default mode, this procedure is required.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree mode {pvst mst rapid-pvst}	Configure a spanning-tree mode. <ul style="list-style-type: none"> • Select pvst to enable PVST+ • Select mst to enable MSTP (and RSTP). For more configuration steps, see Chapter 11 “Configuring MSTP.” • Select rapid-pvst to enable rapid PVST+ (the default setting).

Step	Command	Purpose
3.	interface <i>interface-id</i>	(Recommended for rapid-PVST+ mode only) Enter interface configuration mode, and specify an interface to configure. Valid interfaces include physical ports, VLANs, and port channels. Valid interfaces include physical ports, VLANs, and port channels. The VLAN ID range is 1 to 4094. The port-channel range is 1 to 6.
4.	spanning-tree link-type point-to-point	(Recommended for rapid-PVST+ mode only) Specify that the link type for this port is point-to-point. If you connect this port (local port) to a remote port through a point-to-point link and the local port becomes a designated port, the switch negotiates with the remote port and rapidly transitions the local port to the forwarding state.
5.	end	Return to privileged EXEC mode.
6.	clear spanning-tree detected-protocols	(Recommended for rapid-PVST+ mode only) If any port on the switch is connected to a port on a legacy 802.1D switch, restart the protocol migration process on the entire switch. This step is optional if the designated switch determines that this switch is running rapid PVST+.
7.	show spanning-tree summary and show spanning-tree interface <i>interface-id</i>	Verify your entries.
8.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no spanning-tree mode** global configuration command. To return the port to its default setting, use the **no spanning-tree link-type** interface configuration command.

Disabling Spanning Tree

Spanning tree is enabled by default on VLAN 1 and on all newly created VLANs up to the spanning-tree limit specified in the “Supported Spanning-Tree Instances” section on page 191. Disable spanning tree only if you are sure there are no loops in the network topology.

Caution: When spanning tree is disabled and loops are present in the topology, excessive traffic and indefinite packet duplication can drastically reduce network performance.

Beginning in privileged EXEC mode, follow these steps to disable spanning tree on a per-VLAN basis. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	no spanning-tree vlan <i>vlan-id</i>	Disable spanning tree on a per-VLAN basis. For <i>vlan-id</i> , you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. Valid interfaces include physical ports, VLANs, and port channels. The VLAN ID range is 1 to 4094.
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To re-enable spanning tree, use the **spanning-tree vlan *vlan-id*** global configuration command.

Configuring the Root Switch

The switch maintains a separate spanning-tree instance for each active VLAN configured on it. A bridge ID, consisting of the switch priority and the switch MAC address, is associated with each instance. For each VLAN, the switch with the lowest bridge ID becomes the root switch for that VLAN.

To configure a switch to become the root for the specified VLAN, use the **spanning-tree vlan *vlan-id* root** global configuration command to modify the switch priority from the default value (32768) to a significantly lower value. When you enter this command, the switch checks the switch priority of the root switches for each VLAN. Because of the extended system ID support, the switch sets its own priority for the specified VLAN to 24576 if this value will cause this switch to become the root for the specified VLAN.

If any root switch for the specified VLAN has a switch priority lower than 24576, the switch sets its own priority for the specified VLAN to 4096 less than the lowest switch priority. (4096 is the value of the least-significant bit of a 4-bit switch priority value as shown in Table 36 on page 186.)

Note: The **spanning-tree vlan *vlan-id* root** global configuration command fails if the value necessary to be the root switch is less than 1.

Note: If your network consists of switches that both do and do not support the extended system ID, it is unlikely that the switch with the extended system ID support will become the root switch. The extended system ID increases the switch priority value every time the VLAN number is greater than the priority of the connected switches running older software.

Note: The root switch for each spanning-tree instance should be a backbone or distribution switch. Do not configure an access switch as the spanning-tree primary root.

Use the **diameter** keyword to specify the Layer 2 network diameter (that is, the maximum number of switch hops between any two end stations in the Layer 2 network). When you specify the network diameter, the switch automatically sets an optimal hello time, forward-delay time, and maximum-age time for a network of that diameter, which can significantly reduce the convergence time. You can use the **hello** keyword to override the automatically calculated hello time.

Note: After configuring the switch as the root switch, we recommend that you avoid manually configuring the hello time, forward-delay time, and maximum-age time by

using the **spanning-tree vlan *vlan-id* hello-time**, **spanning-tree vlan *vlan-id* forward-time**, and the **spanning-tree vlan *vlan-id* max-age** global configuration commands.

Beginning in privileged EXEC mode, follow these steps to configure a switch to become the root for the specified VLAN. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree vlan <i>vlan-id</i> root primary [diameter <i>net-diameter</i> [hello-time <i>seconds</i>]]	<p>Configure a switch to become the root for the specified VLAN.</p> <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. <p>Valid interfaces include physical ports, VLANs, and port channels. The VLAN ID range is 1 to 4094.</p> <ul style="list-style-type: none"> (Optional) For diameter <i>net-diameter</i>, specify the maximum number of switches between any two end stations. The range is 2 to 7. (Optional) For hello-time <i>seconds</i>, specify the interval in seconds between the generation of configuration messages by the root switch. The range is 1 to 10 seconds; the default is 2 seconds. <p>Note: When you enter this command without the optional keywords, the switch recalculates the forward-time, hello-time, max-age, and priority settings. If you had previously configured these parameters, the switch recalculates them.</p>
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree detail	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree vlan *vlan-id* root** global configuration command.

Configuring a Secondary Root Switch

When you configure a switch that supports the extended system ID as the secondary root, the switch priority is modified from the default value (32768) to 28672. The switch is then likely to become the root switch for the specified VLAN if the primary root switch fails. This is assuming that the other network switches use the default switch priority of 32768 and therefore are unlikely to become the root switch.

You can execute this command on more than one switch to configure multiple backup root switches. Use the same network diameter and hello-time values as you used when you configured the primary root switch with the **spanning-tree vlan *vlan-id* root primary** global configuration command.

Beginning in privileged EXEC mode, follow these steps to configure a switch to become the secondary root for the specified VLAN. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree vlan <i>vlan-id</i> root secondary [diameter <i>net-diameter</i> [hello-time <i>seconds</i>]]	<p>Configure a switch to become the secondary root for the specified VLAN.</p> <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. Valid interfaces include physical ports, VLANs, and port channels. The VLAN ID range is 1 to 4094. (Optional) For diameter <i>net-diameter</i>, specify the maximum number of switches between any two end stations. The range is 2 to 7. (Optional) For hello-time <i>seconds</i>, specify the interval in seconds between the generation of configuration messages by the root switch. The range is 1 to 10 seconds; the default is 2 seconds. <p>Use the same network diameter and hello-time values that you used when configuring the primary root switch. See the “Configuring the Root Switch” section on page 196.</p>
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree detail	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree vlan *vlan-id* root** global configuration command.

Configuring the Port Priority

If a loop occurs, spanning tree uses the port priority when selecting an interface to put into the forwarding state. You can assign higher priority values (lower numerical values) to interfaces that you want selected first and lower priority values (higher numerical values) that you want selected last. If all interfaces have the same priority value, spanning tree puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

Beginning in privileged EXEC mode, follow these steps to configure the port priority of an interface. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	<p>Enter interface configuration mode, and specify an interface to configure.</p> <p>Valid interfaces include physical interfaces and port-channel logical interfaces (port-channel <i>port-channel-number</i>).</p>

Step	Command	Purpose
3.	spanning-tree port-priority <i>priority</i>	Configure the port priority for an interface. For <i>priority</i> , the range is 0 to 240 in increments of 16. The default is 128. The lower the number, the higher the priority. Valid priority values are 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224, and 240. All other values are rejected.
4.	spanning-tree vlan <i>vlan-id</i> port-priority <i>priority</i>	Configure the VLAN port priority for an interface. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. Valid interfaces include physical ports, VLANs, and port channels. The VLAN ID range is 1 to 4094. For <i>priority</i>, the range is 0 to 240 in increments of 16. The default is 128. The lower the number, the higher the priority. Valid priority values are 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224, and 240. All other values are rejected.
5.	end	Return to privileged EXEC mode.
6.	show spanning-tree interface <i>interface-id</i> or show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Note: The **show spanning-tree interface** *interface-id* privileged EXEC command displays information only if the port is in a link-up operative state. Otherwise, you can use the **show running-config interface** privileged EXEC command to confirm the configuration.

To return the interface to its default setting, use the **no spanning-tree [vlan *vlan-id*] port-priority** interface configuration command. For information on how to configure load sharing on trunk ports by using spanning-tree port priorities, see the “Load Sharing Using STP” section on page 262.

Configuring the Path Cost

The spanning-tree path cost default value is derived from the media speed of an interface. If a loop occurs, spanning tree uses cost when selecting an interface to put in the forwarding state. You can assign lower cost values to interfaces that you want selected first and higher cost values that you want selected last. If all interfaces have the same cost value, spanning tree puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

Beginning in privileged EXEC mode, follow these steps to configure the cost of an interface. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify an interface to configure. Valid interfaces include physical interfaces and port-channel logical interfaces (port-channel <i>port-channel-number</i>).
3.	spanning-tree cost <i>cost</i>	Configure the cost for an interface. If a loop occurs, spanning tree uses the path cost when selecting an interface to place into the forwarding state. A lower path cost represents higher-speed transmission. For <i>cost</i> , the range is 1 to 200000000; the default value is derived from the media speed of the interface. Note: Ports 15 and 16 to the management module have a non-standard default of 100.
4.	spanning-tree vlan <i>vlan-id</i> cost <i>cost</i>	Configure the cost for a VLAN. If a loop occurs, spanning tree uses the path cost when selecting an interface to place into the forwarding state. A lower path cost represents higher-speed transmission. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. Valid interfaces include physical ports, VLANs, and port channels. The VLAN ID range is 1 to 4094. <ul style="list-style-type: none"> For <i>cost</i>, the range is 1 to 200000000; the default value is derived from the media speed of the interface.
5.	end	Return to privileged EXEC mode.
6.	show spanning-tree interface <i>interface-id</i> or show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Note: The **show spanning-tree interface** *interface-id* privileged EXEC command displays information only for ports that are in a link-up operative state. Otherwise, you can use the **show running-config** privileged EXEC command to confirm the configuration.

To return the interface to its default setting, use the **no spanning-tree [vlan *vlan-id*] cost** interface configuration command. For information on how to configure load sharing on trunk ports by using spanning-tree path costs, see the “Load Sharing Using STP” section on page 262.

Configuring the Switch Priority of a VLAN

You can configure the switch priority and make it more likely that the switch will be chosen as the root switch.

Note: Exercise care when using this command. For most situations, we recommend that you use the **spanning-tree vlan *vlan-id* root primary** and the **spanning-tree vlan *vlan-id* root secondary** global configuration commands to modify the switch priority.

Beginning in privileged EXEC mode, follow these steps to configure the switch priority of a VLAN. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree vlan <i>vlan-id</i> priority <i>priority</i>	Configure the switch priority of a VLAN. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. Valid interfaces include physical ports, VLANs, and port channels. The VLAN ID range is 1 to 4094. For <i>priority</i>, the range is 0 to 61440 in increments of 4096; the default is 32768. The lower the number, the more likely the switch will be chosen as the root switch. Valid priority values are 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. All other values are rejected.
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree vlan *vlan-id* priority** global configuration command.

Configuring Spanning-Tree Timers

Table 39 describes the timers that affect the entire spanning-tree performance.

Table 39. Spanning-Tree Timers

Variable	Description
Hello timer	Determines how often the switch broadcasts hello messages to other switches.
Forward-delay timer	Determines how long each of the listening and learning states last before the interface begins forwarding.
Maximum-age timer	Determines the amount of time the switch stores protocol information received on an interface.

The sections that follow provide the configuration steps.

Configuring the Hello Time

You can configure the interval between the generation of configuration messages by the root switch by changing the hello time.

Note: Exercise care when using this command. For most situations, we recommend that you use the **spanning-tree vlan *vlan-id* root primary** and the **spanning-tree vlan *vlan-id* root secondary** global configuration commands to modify the hello time.

Beginning in privileged EXEC mode, follow these steps to configure the hello time of a VLAN. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree vlan <i>vlan-id</i> hello-time <i>seconds</i>	Configure the hello time of a VLAN. The hello time is the interval between the generation of configuration messages by the root switch. These messages mean that the switch is alive. <ul style="list-style-type: none">For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. Valid interfaces include physical ports, VLANs, and port channels. The VLAN ID range is 1 to 4094.For <i>seconds</i>, the range is 1 to 10; the default is 2.
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree vlan *vlan-id* hello-time** global configuration command.

Configuring the Forwarding-Delay Time for a VLAN

Beginning in privileged EXEC mode, follow these steps to configure the forwarding-delay time for a VLAN. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree vlan <i>vlan-id</i> forward-time <i>seconds</i>	Configure the forward time of a VLAN. The forward delay is the number of seconds a port waits before changing from its spanning-tree learning and listening states to the forwarding state. <ul style="list-style-type: none">For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. Valid interfaces include physical ports, VLANs, and port channels. The VLAN ID range is 1 to 4094.For <i>seconds</i>, the range is 4 to 30; the default is 15.
3.	end	Return to privileged EXEC mode.

Step	Command	Purpose
4.	show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree vlan *vlan-id* forward-time** global configuration command.

Configuring the Maximum-Aging Time for a VLAN

Beginning in privileged EXEC mode, follow these steps to configure the maximum-aging time for a VLAN. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree vlan <i>vlan-id</i> max-age <i>seconds</i>	Configure the maximum-aging time of a VLAN. The maximum-aging time is the number of seconds a switch waits without receiving spanning-tree configuration messages before attempting a reconfiguration. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. Valid interfaces include physical ports, VLANs, and port channels. The VLAN ID range is 1 to 4094. For <i>seconds</i>, the range is 6 to 40; the default is 20.
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree vlan *vlan-id* max-age** global configuration command.

Displaying the Spanning-Tree Status

To display the spanning-tree status, use one or more of the privileged EXEC commands in Table 40:

Table 40. Commands for Displaying Spanning-Tree Status

Command	Purpose
show spanning-tree active	Displays spanning-tree information on active interfaces only.
show spanning-tree detail	Displays a detailed summary of interface information.
show spanning-tree interface <i>interface-id</i>	Displays spanning-tree information for the specified interface.
show spanning-tree summary [totals]	Displays a summary of port states or displays the total lines of the STP state section.

You can clear spanning-tree counters by using the **clear spanning-tree [interface *interface-id*]** privileged EXEC command.

For information about other keywords for the **show spanning-tree** privileged EXEC command, refer to the command reference for this release.

Chapter 11. Configuring MSTP

This chapter describes how to configure the Cisco implementation of the IEEE 802.1s Multiple STP (MSTP) on your switch.

The MSTP enables multiple VLANs to be mapped to the same spanning-tree instance, thereby reducing the number of spanning-tree instances needed to support a large number of VLANs. The MSTP provides for multiple forwarding paths for data traffic and enables load balancing. It improves the fault tolerance of the network because a failure in one instance (forwarding path) does not affect other instances (forwarding paths). The most common initial deployment of MSTP is in the backbone and distribution layers of a Layer 2 switched network; this deployment provides the highly-available network required in a service-provider environment.

When the switch is in the multiple spanning-tree (MST) mode, the Rapid Spanning Tree Protocol (RSTP), which is based on IEEE 802.1w, is automatically enabled. The RSTP provides rapid convergence of the spanning tree through explicit handshaking that eliminates the IEEE 802.1D forwarding delay and quickly transitions root ports and designated ports to the forwarding state.

Both MSTP and RSTP improve the spanning-tree operation and maintain backward compatibility with equipment that is based on the (original) 802.1D spanning tree, with existing Cisco-proprietary Multiple Instance STP (MISTP), and with existing Cisco per-VLAN spanning-tree plus (PVST+) and rapid per-VLAN spanning-tree plus (rapid PVST+). For information about PVST+ and rapid PVST+, see Chapter 10 “Configuring STP.” For information about other spanning-tree features such as Port Fast, UplinkFast, root guard, and so forth, see Chapter 12 “Configuring Optional Spanning-Tree Features.”

Note: For complete syntax and usage information for the commands used in this chapter, refer to the *command reference* for this release.

This chapter consists of these sections:

- Understanding MSTP, on page 205
- Understanding RSTP, on page 210
- Configuring MSTP Features, on page 215
- Displaying the MST Configuration and Status, on page 226

Understanding MSTP

MSTP, which uses RSTP for rapid convergence, enables VLANs to be grouped into a spanning-tree instance, with each instance having a spanning-tree topology independent of other spanning-tree instances. This architecture provides multiple forwarding paths for data traffic, enables load balancing, and reduces the number of spanning-tree instances required to support a large number of VLANs.

These sections describe how the MSTP works:

- Multiple Spanning-Tree Regions, on page 206
- IST, CIST, and CST, on page 206
- Hop Count, on page 208
- Boundary Ports, on page 209
- Interoperability with 802.1D STP, on page 209

For configuration information, see the “Configuring MSTP Features” section on page 215.

Multiple Spanning-Tree Regions

For switches to participate in multiple spanning-tree (MST) instances, you must consistently configure the switches with the same MST configuration information. A collection of interconnected switches that have the same MST configuration comprises an MST region as shown in Figure 38. on page 208.

The MST configuration determines to which MST region each switch belongs. The configuration includes the name of the region, the revision number, and the MST VLAN-to-instance assignment map. You configure the switch for a region by using the **spanning-tree mst configuration** global configuration command, after which the switch enters the MST configuration mode. From this mode, you can map VLANs to an MST instance by using the **instance** MST configuration command, specify the region name by using the **name** MST configuration command, and set the revision number by using the **revision** MST configuration command.

A region can have one member or multiple members with the same MST configuration; each member must be capable of processing RSTP bridge protocol data units (BPDUs). There is no limit to the number of MST regions in a network, but each region can support up to 16 spanning-tree instances. You can assign a VLAN to only one spanning-tree instance at a time.

IST, CIST, and CST

Unlike PVST+ and rapid PVST+ in which all the spanning-tree instances are independent, the MSTP establishes and maintains two types of spanning trees:

- An internal spanning tree (IST), which is the spanning tree that runs in an MST region.

Within each MST region, the MSTP maintains multiple spanning-tree instances. Instance 0 is a special instance for a region, known as the internal spanning tree (IST). All other MST instances are numbered from 1 to 15.

The IST is the only spanning-tree instance that sends and receives BPDUs; all of the other spanning-tree instance information is contained in M-records, which are encapsulated within MSTP BPDUs. Because the MSTP BPDU carries information for all instances, the number of BPDUs that need to be processed by a switch to support multiple spanning-tree instances is significantly reduced.

All MST instances within the same region share the same protocol timers, but each MST instance has its own topology parameters, such as root switch ID, root path cost, and so forth. By default, all VLANs are assigned to the IST.

An MST instance is local to the region; for example, MST instance 1 in region A is independent of MST instance 1 in region B, even if regions A and B are interconnected.

- A common and internal spanning tree (CIST), which is a collection of the ISTs in each MST region, and the common spanning tree (CST) that interconnects the MST regions and single spanning trees.

The spanning tree computed in a region appears as a subtree in the CST that encompasses the entire switched domain. The CIST is formed as a result of the spanning-tree algorithm running between switches that support the 802.1w, 802.1s, and 802.1D protocols. The CIST inside an MST region is the same as the CST outside a region.

For more information, see the “Operations Within an MST Region” section on page 207 and the “Operations Between MST Regions” section on page 207.

Operations Within an MST Region

The IST connects all the MSTP switches in a region. When the IST converges, the root of the IST becomes the IST master (shown in Figure 38, on page 208), which is the switch within the region with the lowest bridge ID and path cost to the CST root. The IST master also is the CST root if there is only one region within the network. If the CST root is outside the region, one of the MSTP switches at the boundary of the region is selected as the IST master.

When an MSTP switch initializes, it sends BPDUs claiming itself as the root of the CST and the IST master, with both of the path costs to the CST root and to the IST master set to zero. The switch also initializes all of its MST instances and claims to be the root for all of them. If the switch receives superior MST root information (lower bridge ID, lower path cost, and so forth) than currently stored for the port, it relinquishes its claim as the IST master.

During initialization, a region might have many subregions, each with its own IST master. As switches receive superior IST information, they leave their old subregions and join the new subregion that might contain the true IST master. Thus all subregions shrink, except for the one that contains the true IST master.

For correct operation, all switches in the MST region must agree on the same IST master. Therefore, any two switches in the region synchronize their port roles for an MST instance only if they converge to a common IST master.

Operations Between MST Regions

If there are multiple regions or legacy 802.1D switches within the network, MSTP establishes and maintains the CST, which includes all MST regions and all legacy STP switches in the network. The MST instances combine with the IST at the boundary of the region to become the CST.

The IST connects all the MSTP switches in the region and appears as a subtree in the CST that encompasses the entire switched domain, with the root of the subtree being the IST master. The MST region appears as a virtual switch to adjacent STP switches and MST regions.

Figure 38 shows a network with three MST regions and a legacy 802.1D switch (D). The IST master for region 1 (A) is also the CST root. The IST master for region 2 (B) and the IST master for region 3 (C) are the roots for their respective subtrees within the CST. The RSTP runs in all regions.

Figure 38. MST Regions, IST Masters, and the CST Root

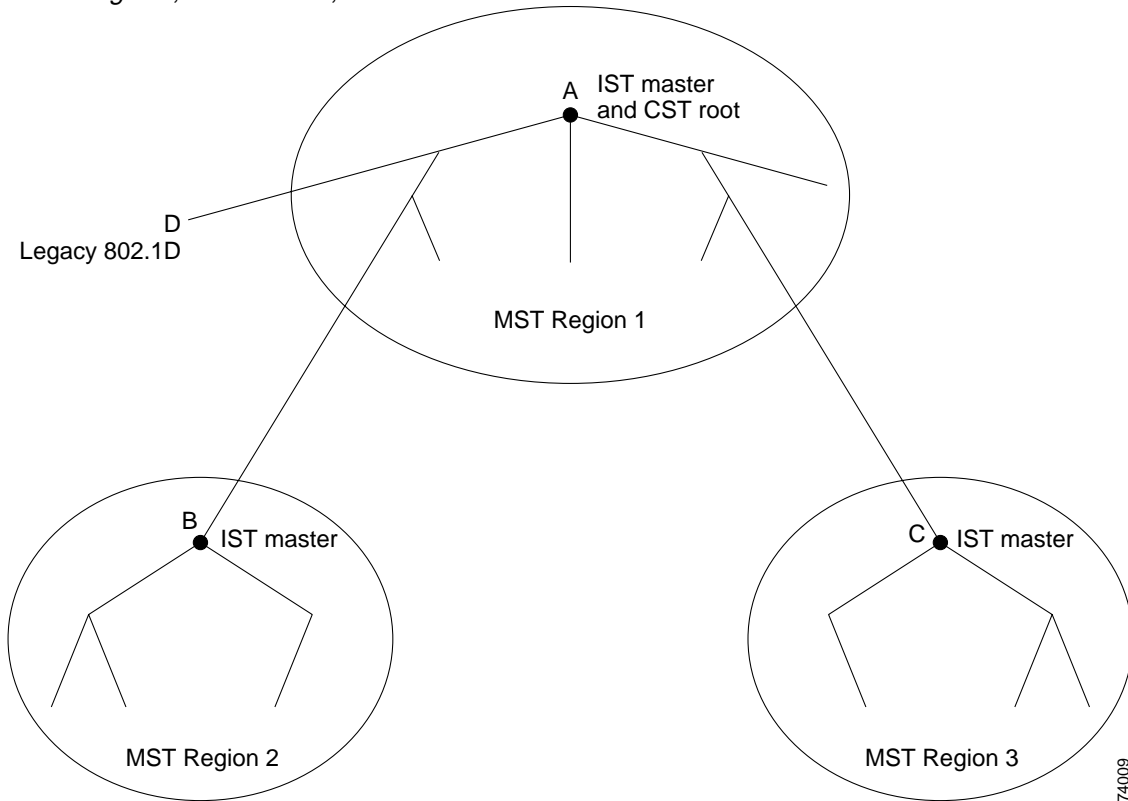


Figure 38 does not show additional MST instances for each region. Note that the topology of MST instances can be different from that of the IST for the same region.

Only the CST instance sends and receives BPDUs, and MST instances add their spanning-tree information into the BPDUs to interact with neighboring switches and compute the final spanning-tree topology. Because of this, the spanning-tree parameters related to BPDU transmission (for example, hello time, forward time, max-age, and max-hops) are configured only on the CST instance but affect all MST instances. Parameters related to the spanning-tree topology (for example, switch priority, port VLAN cost, port VLAN priority) can be configured on both the CST instance and the MST instance.

MSTP switches use version 3 RSTP BPDUs or 802.1D STP BPDUs to communicate with legacy 802.1D switches. MSTP switches use MSTP BPDUs to communicate with MSTP switches.

Hop Count

The IST and MST instances do not use the message-age and maximum-age information in the configuration BPDU to compute the spanning-tree topology. Instead, they use the path cost to the root and a hop-count mechanism similar to the IP time-to-live (TTL) mechanism.

By using the **spanning-tree mst max-hops** global configuration command, you can configure the maximum hops inside the region and apply it to the IST and all MST instances in that region. The hop count achieves the same result as the message-age information (determines when to trigger a reconfiguration). The root switch of the instance always sends a BPDU (or M-record) with a cost of 0 and the hop count set to the maximum value. When a switch receives this BPDU, it decrements the received remaining hop count by one and propagates this value as the remaining hop count in

the BPDUs it generates. When the count reaches zero, the switch discards the BPDU and ages the information held for the port.

The message-age and maximum-age information in the RSTP portion of the BPDU remain the same throughout the region, and the same values are propagated by the region's designated ports at the boundary.

Boundary Ports

A boundary port is a port that connects an MST region to a single spanning-tree region running RSTP, to a single spanning-tree region running PVST+ or rapid PVST+, or to another MST region with a different MST configuration. A boundary port also connects to a LAN, the designated switch of which is either a single spanning-tree switch or a switch with a different MST configuration.

At the boundary, the roles of the MST ports do not matter, and their state is forced to be the same as the IST port state (MST ports at the boundary are in the forwarding state only when the IST port is forwarding). An IST port at the boundary can have any port role except a backup port role.

On a shared boundary link, the MST ports wait in the blocking state for the forward-delay time to expire before transitioning to the learning state. The MST ports wait another forward-delay time before transitioning to the forwarding state.

If the boundary port is on a point-to-point link and it is the IST root port, the MST ports transition to the forwarding state as soon as the IST port transitions to the forwarding state.

If the IST port is a designated port on a point-to-point link and if the IST port transitions to the forwarding state because of an agreement received from its peer port, the MST ports also immediately transition to the forwarding state.

If a boundary port transitions to the forwarding state in an IST instance, it is forwarding in all MST instances, and a topology change is triggered. If a boundary port with the IST root or designated port role receives a topology change notice external to the MST cloud, the MSTP switch triggers a topology change in the IST instance and in all the MST instances active on that port.

Interoperability with 802.1D STP

A switch running MSTP supports a built-in protocol migration mechanism that enables it to interoperate with legacy 802.1D switches. If this switch receives a legacy 802.1D configuration BPDU (a BPDU with the protocol version set to 0), it sends only 802.1D BPDUs on that port. An MSTP switch can also detect that a port is at the boundary of a region when it receives a legacy BPDU, an MSTP BPDU (version 3) associated with a different region, or an RSTP BPDU (version 2).

However, the switch does not automatically revert to the MSTP mode if it no longer receives 802.1D BPDUs because it cannot determine whether the legacy switch has been removed from the link unless the legacy switch is the designated switch. Also, a switch might continue to assign a boundary role to a port when the switch to which this switch is connected has joined the region. To restart the protocol migration process (force the renegotiation with neighboring switches), use the **clear spanning-tree detected-protocols** privileged EXEC command.

If all the legacy switches on the link are RSTP switches, they can process MSTP BPDUs as if they are RSTP BPDUs. Therefore, MSTP switches send either a version 0 configuration and TCN BPDUs or version 3 MSTP BPDUs on a boundary port. A

boundary port connects to a LAN, the designated switch of which is either a single spanning-tree switch or a switch with a different MST configuration.

Understanding RSTP

The RSTP takes advantage of point-to-point wiring and provides rapid convergence of the spanning tree. Reconfiguration of the spanning tree can occur in less than 1 second (in contrast to 50 seconds with the default settings in the 802.1D spanning tree), which is critical for networks carrying delay-sensitive traffic such as voice and video.

This section describes how the RSTP works:

- Port Roles and the Active Topology, on page 210
- Rapid Convergence, on page 211
- Synchronization of Port Roles, on page 212
- Bridge Protocol Data Unit Format and Processing, on page 213
- “Topology Changes” section on page 214

For configuration information, see the “Configuring MSTP Features” section on page 215.

Port Roles and the Active Topology

The RSTP provides rapid convergence of the spanning tree by assigning port roles and by determining the active topology. The RSTP builds upon the IEEE 802.1D STP to select the switch with the highest switch priority (lowest numerical priority value) as the root switch as described in the “Spanning-Tree Topology and BPDUs” section on page 184. Then the RSTP assigns one of these port roles to individual ports:

- Root port—Provides the best path (lowest cost) when the switch forwards packets to the root switch.
- Designated port—Connects to the designated switch, which incurs the lowest path cost when forwarding packets from that LAN to the root switch. The port through which the designated switch is attached to the LAN is called the designated port.
- Alternate port—Offers an alternate path toward the root switch to that provided by the current root port.
- Backup port—Acts as a backup for the path provided by a designated port toward the leaves of the spanning tree. A backup port can exist only when two ports are connected together in a loopback by a point-to-point link or when a switch has two or more connections to a shared LAN segment.
- Disabled port—Has no role within the operation of the spanning tree.

A port with the root or a designated port role is included in the active topology. A port with the alternate or backup port role is excluded from the active topology.

In a stable topology with consistent port roles throughout the network, the RSTP ensures that every root port and designated port immediately transition to the forwarding state while all alternate and backup ports are always in the discarding state (equivalent to blocking in 802.1D). The port state controls the operation of the forwarding and learning processes. Table 41 provides a comparison of 802.1D and RSTP port states.

Table 41. Port State Comparison

Operational Status	STP Port State (802.1D)	RSTP Port State	Is Port Included in the Active Topology?
Enabled	Blocking	Discarding	No
Enabled	Listening	Discarding	No
Enabled	Learning	Learning	Yes
Enabled	Forwarding	Forwarding	Yes
Disabled	Disabled	Discarding	No

To be consistent with Cisco STP implementations, this guide documents the port state as *blocking* instead of *discarding*. Designated ports start in the listening state.

Rapid Convergence

The RSTP provides for rapid recovery of connectivity following the failure of a switch, a switch port, or a LAN. It provides rapid convergence for edge ports, new root ports, and ports connected through point-to-point links as follows:

- Edge ports—If you configure a port as an edge port on an RSTP switch by using the **spanning-tree portfast** interface configuration command, the edge port immediately transitions to the forwarding state. An edge port is the same as a Port Fast-enabled port, and you should enable it only on ports that connect to a single end station.
- Root ports—If the RSTP selects a new root port, it blocks the old root port and immediately transitions the new root port to the forwarding state.
- Point-to-point links—If you connect a port to another port through a point-to-point link and the local port becomes a designated port, it negotiates a rapid transition with the other port by using the proposal-agreement handshake to ensure a loop-free topology.

As shown in Figure 39, Switch A is connected to Switch B through a point-to-point link, and all of the ports are in the blocking state. Assume that the priority of Switch A is a smaller numerical value than the priority of Switch B. Switch A sends a proposal message (a configuration BPDU) with the proposal flag set to Switch B, proposing itself as the designated switch.

After receiving the proposal message, Switch B selects as its new root port the port from which the proposal message was received, forces all nonedge ports to the blocking state, and sends an agreement message (a BPDU with the agreement flag set) through its new root port.

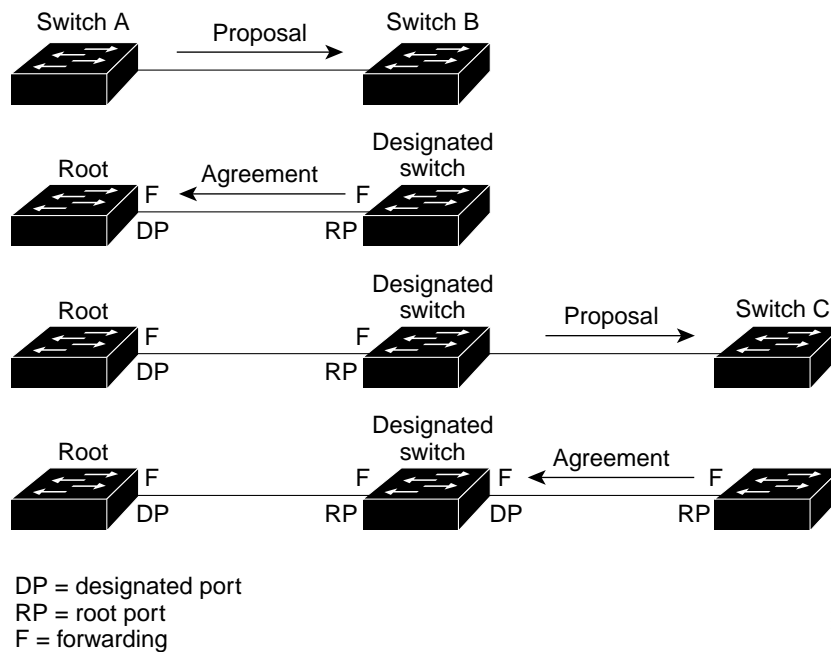
After receiving Switch B's agreement message, Switch A also immediately transitions its designated port to the forwarding state. No loops in the network are formed because Switch B blocked all of its nonedge ports and because there is a point-to-point link between Switches A and B.

When Switch C is connected to Switch B, a similar set of handshaking messages are exchanged. Switch C selects the port connected to Switch B as its root port, and both ends immediately transition to the forwarding state. With each iteration of this handshaking process, one more switch joins the active topology. As the network converges, this proposal-agreement handshaking progresses from the root toward the leaves of the spanning tree.

The switch determines the link type from the port duplex mode: a full-duplex port is considered to have a point-to-point connection; a half-duplex port is considered to have a shared connection. You can override the default setting that is

determined by the duplex setting by using the **spanning-tree link-type** interface configuration command.

Figure 39. Proposal and Agreement Handshaking for Rapid Convergence



Synchronization of Port Roles

When the switch receives a proposal message on one of its ports and that port is selected as the new root port, the RSTP forces all other ports to synchronize with the new root information.


The switch is synchronized with superior root information received on the root port if all other ports are synchronized. An individual port on the switch is synchronized if

- That port is in the blocking state
- It is an edge port (a port configured to be at the edge of the network)

If a designated port is in the forwarding state and is not configured as an edge port, it transitions to the blocking state when the RSTP forces it to synchronize with new root information. In general, when the RSTP forces a port to synchronize with root information and the port does not satisfy any of the above conditions, its port state is set to blocking.

After ensuring all of the ports are synchronized, the switch sends an agreement message to the designated switch corresponding to its root port. When the switches connected by a point-to-point link are in agreement about their port roles, the RSTP immediately transitions the port states to forwarding. The sequence of events is shown in Figure 40.

Figure 40. Sequence of Events During Rapid Convergence



Bridge Protocol Data Unit Format and Processing

The RSTP BPDUs are the same as the IEEE 802.1D BPDUs except that the protocol version is set to 2. A new one-byte version 1 Length field is set to zero, which means that no version 1 protocol information is present. Table 42 shows the RSTP flag fields.

Table 42. RSTP BPDUs Flags

Bit	Function
0	Topology change (TC)
1	Proposal
2–3:	Port role:
00	Unknown
01	Alternate port
10	Root port
11	Designated port
4	Learning
5	Forwarding
6	Agreement
7	Topology change acknowledgement (TCA)

The sending switch sets the proposal flag in the RSTP BPDUs to propose itself as the designated switch on that LAN. The port role in the proposal message is always set to the designated port.

The sending switch sets the agreement flag in the RSTP BPDUs to accept the previous proposal. The port role in the agreement message is always set to the root port.

The RSTP does not have a separate topology change notification (TCN) BPDUs. It uses the topology change (TC) flag to show the topology changes. However, for

interoperability with 802.1D switches, the RSTP switch processes and generates TCN BPDUs.

The learning and forwarding flags are set according to the state of the sending port.

Processing Superior BPDU Information

If a port receives superior root information (lower bridge ID, lower path cost, and so forth) than currently stored for the port, the RSTP triggers a reconfiguration. If the port is proposed and is selected as the new root port, RSTP forces all the other ports to synchronize.

If the BPDU received is an RSTP BPDU with the proposal flag set, the switch sends an agreement message after all of the other ports are synchronized. If the BPDU is an 802.1D BPDU, the switch does not set the proposal flag and starts the forward-delay timer for the port. The new root port requires twice the forward-delay time to transition to the forwarding state.

If the superior information received on the port causes the port to become a backup or alternate port, RSTP sets the port to the blocking state but does not send the agreement message. The designated port continues sending BPDUs with the proposal flag set until the forward-delay timer expires, at which time the port transitions to the forwarding state.

Processing Inferior BPDU Information

If a designated port receives an inferior BPDU (higher bridge ID, higher path cost, and so forth than currently stored for the port) with a designated port role, it immediately replies with its own information.

Topology Changes

This section describes the differences between the RSTP and the 802.1D in handling spanning-tree topology changes.

- **Detection**—Unlike 802.1D in which *any* transition between the blocking and the forwarding state causes a topology change, *only* transitions from the blocking to the forwarding state cause a topology change with RSTP (only an increase in connectivity is considered a topology change). State changes on an edge port do not cause a topology change. When an RSTP switch detects a topology change, it flushes the learned information on all of its nonedge ports except on those from which it received the TC notification.
- **Notification**—Unlike 802.1D, which uses TCN BPDUs, the RSTP does not use them. However, for 802.1D interoperability, an RSTP switch processes and generates TCN BPDUs.
- **Acknowledgement**—When an RSTP switch receives a TCN message on a designated port from an 802.1D switch, it replies with an 802.1D configuration BPDU with the TCA bit set. However, if the TC-while timer (the same as the topology-change timer in 802.1D) is active on a root port connected to an 802.1D switch and a configuration BPDU with the TCA bit set is received, the TC-while timer is reset.

This behavior is only required to support 802.1D switches. The RSTP BPDUs never have the TCA bit set.

- **Propagation**—When an RSTP switch receives a TC message from another switch through a designated or root port, it propagates the change to all of its nonedge, designated ports and to the root port (excluding the port on which it is received). The switch starts the TC-while timer for all such ports and flushes the information learned on them.

- Protocol migration—For backward compatibility with 802.1D switches, RSTP selectively sends 802.1D configuration BPDUs and TCN BPDUs on a per-port basis.

When a port is initialized, the migrate-delay timer is started (specifies the minimum time during which RSTP BPDUs are sent), and RSTP BPDUs are sent. While this timer is active, the switch processes all BPDUs received on that port and ignores the protocol type.

If the switch receives an 802.1D BPDU after the port's migration-delay timer has expired, it assumes that it is connected to an 802.1D switch and starts using only 802.1D BPDUs. However, if the RSTP switch is using 802.1D BPDUs on a port and receives an RSTP BPDU after the timer has expired, it restarts the timer and starts using RSTP BPDUs on that port.

Configuring MSTP Features

These sections describe how to configure basic MSTP features:

- Default MSTP Configuration, on page 215
- MSTP Configuration Guidelines, on page 216
- Specifying the MST Region Configuration and Enabling MSTP, on page 217 (required)
- Configuring the Root Switch, on page 218 (optional)
- Configuring a Secondary Root Switch, on page 219 (optional)
- Configuring the Port Priority, on page 220 (optional)
- Configuring the Path Cost, on page 221 (optional)
- Configuring the Switch Priority, on page 222 (optional)
- Configuring the Hello Time, on page 223 (optional)
- Configuring the Forwarding-Delay Time, on page 224 (optional)
- Configuring the Maximum-Aging Time, on page 224 (optional)
- Configuring the Maximum-Hop Count, on page 224 (optional)
- Specifying the Link Type to Ensure Rapid Transitions, on page 225 (optional)
- Restarting the Protocol Migration Process, on page 225 (optional)

Default MSTP Configuration

Table 43 shows the default MSTP configuration.

Table 43. Default MSTP Configuration

Feature	Default Setting
Spanning-tree mode	PVST+ (Rapid PVST+ and MSTP are disabled)
Switch priority (configurable on a per-CIST interface basis)	32768
Spanning-tree port priority (configurable on a per-CIST interface basis)	128

Table 43. Default MSTP Configuration (continued)

Feature	Default Setting
Spanning-tree port cost (configurable on a per-CIST interface basis)	1000 Mbps: 4 100 Mbps: 100 (for the internal 100 Mbps management module ports) 100 Mbps: 19 (for the external 10/100/1000 Mbps ports) 10 Mbps: 100.
Hello time	2 seconds
Forward-delay time	15 seconds
Maximum-aging time	20 seconds
Maximum hop count	20 hops

For information about the supported number of spanning-tree instances, see the “Supported Spanning-Tree Instances” section on page 191.

MSTP Configuration Guidelines

These are the configuration guidelines for MSTP:

- When you enable MST by using the **spanning-tree mode mst** global configuration command, RSTP is automatically enabled.
- For two or more switches to be in the same MST region, they must have the same VLAN-to-instance map, the same configuration revision number, and the same name.
- The switch supports up to 16 MST instances. The number of VLANs that can be mapped to a particular MST instance is unlimited.
- The UplinkFast and BackboneFast features are not supported with the MSTP.
- PVST+, rapid PVST+, and MSTP are supported, but only one version can be active at any time. (For example, all VLANs run PVST+, all VLANs run rapid PVST+, or all VLANs run MSTP.) For more information, see the “Spanning-Tree Interoperability and Backward Compatibility” section on page 191.
- VTP propagation of the MST configuration is not supported. However, you can manually configure the MST configuration (region name, revision number, and VLAN-to-instance mapping) on each switch within the MST region by using the command-line interface (CLI) or through the SNMP support.
- For load balancing across redundant paths in the network to work, all VLAN-to-instance mapping assignments must match; otherwise, all traffic flows on a single link.
- All MST boundary ports must be forwarding for load balancing between a PVST+ and an MST cloud or between a rapid-PVST+ and an MST cloud. For this to occur, the IST master of the MST cloud should also be the root of the CST. If the MST cloud consists of multiple MST regions, one of the MST regions must contain the CST root, and all of the other MST regions must have a better path to the root contained within the MST cloud than a path through the PVST+ or rapid-PVST+ cloud. You might have to manually configure the switches in the clouds.
- Partitioning the network into a large number of regions is not recommended. However, if this situation is unavoidable, we recommend that you partition the switched LAN into smaller LANs interconnected by routers or non-Layer 2 devices.

Specifying the MST Region Configuration and Enabling MSTP

For two or more switches to be in the same MST region, they must have the same VLAN-to-instance mapping, the same configuration revision number, and the same name.

A region can have one member or multiple members with the same MST configuration; each member must be capable of processing RSTP BPDUs. There is no limit to the number of MST regions in a network, but each region can support up to 16 spanning-tree instances. You can assign a VLAN to only one spanning-tree instance at a time.

Beginning in privileged EXEC mode, follow these steps to specify the MST region configuration and enable MSTP. This procedure is required.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree mst configuration	Enter MST configuration mode.
3.	instance <i>instance-id</i> vlan <i>vlan-range</i>	<p>Map VLANs to an MST instance.</p> <ul style="list-style-type: none"> For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 1 to 15. For vlan <i>vlan-range</i>, the range is 1 to 4094. <p>When you map VLANs to an MST instance, the mapping is incremental, and the VLANs specified in the command are added to or removed from the VLANs that were previously mapped.</p> <p>To specify a VLAN range, use a hyphen; for example, instance 1 vlan 1-63 maps VLANs 1 through 63 to MST instance 1.</p> <p>To specify a VLAN series, use a comma; for example, instance 1 vlan 10, 20, 30 maps VLANs 10, 20, and 30 to MST instance 1.</p>
4.	name <i>name</i>	Specify the configuration name. The <i>name</i> string has a maximum length of 32 characters and is case sensitive.
5.	revision <i>version</i>	Specify the configuration revision number. The range is 0 to 65535.
6.	show pending	Verify your configuration by displaying the pending configuration.
7.	exit	Apply all changes, and return to global configuration mode.
8.	spanning-tree mode mst	<p>Enable MSTP. RSTP is also enabled.</p> <p>Caution: Changing spanning-tree modes can disrupt traffic because all spanning-tree instances are stopped for the previous mode and restarted in the new mode.</p> <p>You cannot run both MSTP and PVST+ or both MSTP and rapid PVST+ at the same time.</p>
9.	end	Return to privileged EXEC mode.
10.	show running-config	Verify your entries.
11.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default MST region configuration, use the **no spanning-tree mst configuration** global configuration command. To return to the default VLAN-to-instance map, use the **no instance** *instance-id* [**vlan** *vlan-range*] MST configuration command. To return to the default name, use the **no name** MST configuration command. To return to the default revision number, use the **no revision** MST configuration command. To re-enable PVST+, use the **no spanning-tree mode** or the **spanning-tree mode pvst** global configuration command.

This example shows how to enter MST configuration mode, map VLANs 10 to 20 to MST instance 1, name the region *region1*, set the configuration revision to 1, display the pending configuration, apply the changes, and return to global configuration mode:

```
Switch(config)# spanning-tree mst configuration
Switch(config-mst)# instance 1 vlan 10-20
Switch(config-mst)# name region1
Switch(config-mst)# revision 1
Switch(config-mst)# show pending
Pending MST configuration
Name      [region1]
Revision  1
Instance  Vlans Mapped
-----
0         1-9,21-4094
1         10-20
-----

Switch(config-mst)# exit
Switch(config)#
```

Configuring the Root Switch

The switch maintains a spanning-tree instance for the group of VLANs mapped to it. A bridge ID, consisting of the switch priority and the switch MAC address, is associated with each instance. The switch with the lowest bridge ID becomes the root switch for the group of VLANs.

To configure a switch to become the root, use the **spanning-tree mst** *instance-id* **root** global configuration command to modify the switch priority from the default value (32768) to a significantly lower value so that the switch becomes the root switch for the specified spanning-tree instance. When you enter this command, the switch checks the switch priorities of the root switches. Because of the extended system ID support, the switch sets its own priority for the specified instance to 24576 if this value will cause this switch to become the root for the specified spanning-tree instance.

If any root switch for the specified instance has a switch priority lower than 24576, the switch sets its own priority to 4096 less than the lowest switch priority. (4096 is the value of the least-significant bit of a 4-bit switch priority value as shown in Table 36 on page 186.)

Note: If your network consists of switches that both do and do not support the extended system ID, it is unlikely that the switch with the extended system ID support will become the root switch. The extended system ID increases the switch priority value every time the VLAN number is greater than the priority of the connected switches running older software.

Note: The root switch for each spanning-tree instance should be a backbone or distribution switch. Do not configure an access switch as the spanning-tree primary root.

Use the **diameter** keyword, which is available only for MST instance 0, to specify the Layer 2 network diameter (that is, the maximum number of switch hops between any

two end stations in the Layer 2 network). When you specify the network diameter, the switch automatically sets an optimal hello time, forward-delay time, and maximum-age time for a network of that diameter, which can significantly reduce the convergence time. You can use the **hello** keyword to override the automatically calculated hello time.

Note: After configuring the switch as the root switch, we recommend that you avoid manually configuring the hello time, forward-delay time, and maximum-age time by using the **spanning-tree mst hello-time**, **spanning-tree mst forward-time**, and the **spanning-tree mst max-age** global configuration commands after configuring the switch as the root switch.

Beginning in privileged EXEC mode, follow these steps to configure a switch as the root switch. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree mst <i>instance-id</i> root primary [diameter <i>net-diameter</i> [hello-time <i>seconds</i>]]	Configure a switch as the root switch. <ul style="list-style-type: none"> For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 15. (Optional) For diameter <i>net-diameter</i>, specify the maximum number of switches between any two end stations. The range is 2 to 7. This keyword is available only for MST instance 0. (Optional) For hello-time <i>seconds</i>, specify the interval in seconds between the generation of configuration messages by the root switch. The range is 1 to 10 seconds; the default is 2 seconds.
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree mst <i>instance-id</i>	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree mst *instance-id* root** global configuration command.

Configuring a Secondary Root Switch

When you configure a switch that supports the extended system ID as the secondary root, the spanning-tree switch priority is modified from the default value (32768) to 28672. The switch is then likely to become the root switch for the specified instance if the primary root switch fails. This is assuming that the other network switches use the default switch priority of 32768 and therefore are unlikely to become the root switch.

You can execute this command on more than one switch to configure multiple backup root switches. Use the same network diameter and hello-time values that you used when you configured the primary root switch with the **spanning-tree mst *instance-id* root primary** global configuration command.

Beginning in privileged EXEC mode, follow these steps to configure a switch as the secondary root switch. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree mst <i>instance-id</i> root secondary [diameter <i>net-diameter</i> [hello-time <i>seconds</i>]]	<p>Configure a switch as the secondary root switch.</p> <ul style="list-style-type: none"> For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 15. (Optional) For diameter <i>net-diameter</i>, specify the maximum number of switches between any two end stations. The range is 2 to 7. This keyword is available only for MST instance 0. (Optional) For hello-time <i>seconds</i>, specify the interval in seconds between the generation of configuration messages by the root switch. The range is 1 to 10 seconds; the default is 2 seconds. <p>Use the same network diameter and hello-time values that you used when configuring the primary root switch. See the “Configuring the Root Switch” section on page 218.</p>
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree mst <i>instance-id</i>	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree mst** *instance-id* **root** global configuration command.

Configuring the Port Priority

If a loop occurs, the MSTP uses the port priority when selecting an interface to put into the forwarding state. You can assign higher priority values (lower numerical values) to interfaces that you want selected first and lower priority values (higher numerical values) that you want selected last. If all interfaces have the same priority value, the MSTP puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

Beginning in privileged EXEC mode, follow these steps to configure the MSTP port priority of an interface. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	<p>Enter interface configuration mode, and specify an interface to configure.</p> <p>Valid interfaces include physical ports and port channels. Valid port-channel numbers are 1 to 6.</p>

Step	Command	Purpose
3.	spanning-tree mst <i>instance-id</i> port-priority <i>priority</i>	Configure the port priority for an MST instance. <ul style="list-style-type: none"> For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 15. For <i>priority</i>, the range is 0 to 240 in increments of 16. The default is 128. The lower the number, the higher the priority. Valid priority values are 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224, and 240. All other values are rejected.
4.	end	Return to privileged EXEC mode.
5.	show spanning-tree mst interface <i>interface-id</i> or show spanning-tree mst <i>instance-id</i>	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Note: The **show spanning-tree mst interface** *interface-id* privileged EXEC command displays information only if the port is in a link-up operative state. Otherwise, you can use the **show running-config interface** privileged EXEC command to confirm the configuration.

To return the interface to its default setting, use the **no spanning-tree mst** *instance-id* **port-priority** interface configuration command.

Configuring the Path Cost

The MSTP path cost default value is derived from the media speed of an interface. If a loop occurs, the MSTP uses cost when selecting an interface to put in the forwarding state. You can assign lower cost values to interfaces that you want selected first and higher cost values that you want selected last. If all interfaces have the same cost value, the MSTP puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

Beginning in privileged EXEC mode, follow these steps to configure the MSTP cost of an interface. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify an interface to configure. Valid interfaces include physical ports and port channels. Valid port-channel numbers are 1 to 6.

Step	Command	Purpose
3.	spanning-tree mst <i>instance-id</i> cost <i>cost</i>	Configure the cost for an MST instance. If a loop occurs, the MSTP uses the path cost when selecting an interface to place into the forwarding state. A lower path cost represents higher-speed transmission. <ul style="list-style-type: none"> For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 15. For <i>cost</i>, the range is 1 to 200000000; the default value is derived from the media speed of the interface.
4.	end	Return to privileged EXEC mode.
5.	show spanning-tree mst interface <i>interface-id</i> or show spanning-tree mst <i>instance-id</i>	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Note: The **show spanning-tree mst interface** *interface-id* privileged EXEC command displays information only for ports that are in a link-up operative state. Otherwise, you can use the **show running-config** privileged EXEC command to confirm the configuration.

To return the interface to its default setting, use the **no spanning-tree mst** *instance-id* **cost** interface configuration command.

Configuring the Switch Priority

You can configure the switch priority and make it more likely that the switch will be chosen as the root switch.

Note: Exercise care when using this command. For most situations, we recommend that you use the **spanning-tree mst** *instance-id* **root primary** and the **spanning-tree mst** *instance-id* **root secondary** global configuration commands to modify the switch priority.

Beginning in privileged EXEC mode, follow these steps to configure the switch priority. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree mst <i>instance-id</i> priority <i>priority</i>	Configure the switch priority for an MST instance. <ul style="list-style-type: none"> For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 15. For <i>priority</i>, the range is 0 to 61440 in increments of 4096; the default is 32768. The lower the number, the more likely the switch will be chosen as the root switch. Valid priority values are 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. All other values are rejected.
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree mst <i>instance-id</i>	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree mst *instance-id* priority** global configuration command.

Configuring the Hello Time

You can configure the interval between the generation of configuration messages by the root switch by changing the hello time.

Note: Exercise care when using this command. For most situations, we recommend that you use the **spanning-tree mst *instance-id* root primary** and the **spanning-tree mst *instance-id* root secondary** global configuration commands to modify the hello time.

Beginning in privileged EXEC mode, follow these steps to configure the hello time for all MST instances. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree mst hello-time <i>seconds</i>	Configure the hello time for all MST instances. The hello time is the interval between the generation of configuration messages by the root switch. These messages mean that the switch is alive. For <i>seconds</i> , the range is 1 to 10; the default is 2.
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree mst	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree mst hello-time** global configuration command.

Configuring the Forwarding-Delay Time

Beginning in privileged EXEC mode, follow these steps to configure the forwarding-delay time for all MST instances. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree mst forward-time <i>seconds</i>	Configure the forward time for all MST instances. The forward delay is the number of seconds a port waits before changing from its spanning-tree learning and listening states to the forwarding state. For <i>seconds</i> , the range is 4 to 30; the default is 15.
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree mst	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree mst forward-time** global configuration command.

Configuring the Maximum-Aging Time

Beginning in privileged EXEC mode, follow these steps to configure the maximum-aging time for all MST instances. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree mst max-age <i>seconds</i>	Configure the maximum-aging time for all MST instances. The maximum-aging time is the number of seconds a switch waits without receiving spanning-tree configuration messages before attempting a reconfiguration. For <i>seconds</i> , the range is 6 to 40; the default is 20.
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree mst	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree mst max-age** global configuration command.

Configuring the Maximum-Hop Count

Beginning in privileged EXEC mode, follow these steps to configure the maximum-hop count for all MST instances. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree mst max-hops <i>hop-count</i>	Specify the number of hops in a region before the BPDU is discarded, and the information held for a port is aged. For <i>hop-count</i> , the range is 1 to 40; the default is 20.

Step	Command	Purpose
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree mst	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree mst max-hops** global configuration command.

Specifying the Link Type to Ensure Rapid Transitions

If you connect a port to another port through a point-to-point link and the local port becomes a designated port, the RSTP negotiates a rapid transition with the other port by using the proposal-agreement handshake to ensure a loop-free topology as described in the “Rapid Convergence” section on page 211.

By default, the link type is determined from the duplex mode of the interface: a full-duplex port is considered to have a point-to-point connection; a half-duplex port is considered to have a shared connection. If you have a half-duplex link physically connected point-to-point to a single port on a remote switch running MSTP, you can override the default setting of the link type and enable rapid transitions to the forwarding state.

Beginning in privileged EXEC mode, follow these steps to override the default link-type setting. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
1.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to configure. Valid interfaces include physical ports, VLANs, and port channels. Valid VLAN IDs are 1 to 4094; valid port-channel numbers are 1 to 6.
2.	spanning-tree link-type point-to-point	Specify that the link type of a port is point-to-point.
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree mst interface <i>interface-id</i>	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree link-type** interface configuration command.

Restarting the Protocol Migration Process

A switch running MSTP supports a built-in protocol migration mechanism that enables it to interoperate with legacy 802.1D switches. If this switch receives a legacy 802.1D configuration BPDU (a BPDU with the protocol version set to 0), it sends only 802.1D BPDUs on that port. An MSTP switch can also detect that a port is at the boundary of a region when it receives a legacy BPDU, an MST BPDU (version 3) associated with a different region, or an RST BPDU (version 2).

However, the switch does not automatically revert to the MSTP mode if it no longer receives 802.1D BPDUs because it cannot determine whether the legacy switch has been removed from the link unless the legacy switch is the designated switch. A switch also might continue to assign a boundary role to a port when the switch to which it is connected has joined the region.

To restart the protocol migration process (force the renegotiation with neighboring switches) on the switch, use the **clear spanning-tree detected-protocols** privileged EXEC command.

To restart the protocol migration process on a specific interface, use the **clear spanning-tree detected-protocols interface *interface-id*** privileged EXEC command.

Displaying the MST Configuration and Status

To display the spanning-tree status, use one or more of the privileged EXEC commands in Table 44:

Table 44. Commands for Displaying MST Status

Command	Purpose
show spanning-tree mst configuration	Displays the MST region configuration.
show spanning-tree mst <i>instance-id</i>	Displays MST information for the specified instance.
show spanning-tree mst interface <i>interface-id</i>	Displays MST information for the specified interface. Valid interfaces include physical ports, VLANs, and port channels. The valid VLAN range is 1 to 4094; the valid port-channel range is 1 to 6.

For information about other keywords for the **show spanning-tree** privileged EXEC command, refer to the *command reference* for this release.

Chapter 12. Configuring Optional Spanning-Tree Features

This chapter describes how to configure optional spanning-tree features on your switch. You can configure all of these features when your switch is running the per-VLAN spanning-tree plus (PVST+). You can configure only the noted features when your switch is running the Multiple Spanning Tree Protocol (MSTP) or the rapid per-VLAN spanning-tree plus (rapid-PVST+) protocol.

For information on configuring the PVST+ and rapid PVST+, see Chapter 10 “Configuring STP.” For information about the Multiple Spanning Tree Protocol (MSTP) and how to map multiple VLANs to the same spanning-tree instance, see Chapter 11 “Configuring MSTP.”

Note: For complete syntax and usage information for the commands used in this chapter, refer to the *command reference* for this release.

This chapter consists of these sections:

- Understanding Optional Spanning-Tree Features, on page 227
- Configuring Optional Spanning-Tree Features, on page 234
- Displaying the Spanning-Tree Status, on page 241

Understanding Optional Spanning-Tree Features

These sections describe how the optional spanning-tree features work:

- Understanding Port Fast, on page 227
- Understanding BPDU Guard, on page 228
- Understanding BPDU Filtering, on page 228
- Understanding UplinkFast, on page 229
- Understanding BackboneFast, on page 231
- Understanding EtherChannel Guard, on page 233
- Understanding Root Guard, on page 233
- Understanding Loop Guard, on page 234

Understanding Port Fast


Port Fast immediately brings an interface configured as an access or trunk port to the forwarding state from a blocking state, bypassing the listening and learning states. You can use Port Fast on ports connected to a single workstation or server, as shown in Figure 41, to allow those devices to immediately connect to the network, rather than waiting for the spanning tree to converge.

Ports connected to a single workstation or server should not receive bridge protocol data units (BPDUs). A port with Port Fast enabled goes through the normal cycle of spanning-tree status changes when the switch is restarted.

Note: Because the purpose of Port Fast is to minimize the time ports must wait for spanning-tree to converge, it is effective only when used on ports connected to end stations. If you enable Port Fast on a port connecting to another switch, you risk creating a spanning-tree loop.

If your switch is running PVST+, rapid PVST+, or MSTP, you can enable this feature by using the **spanning-tree portfast** interface configuration or the **spanning-tree portfast default** global configuration command.

Figure 41. Port Fast-Enabled Ports



Understanding BPDU Guard

The BPDU guard feature can be globally enabled on the switch or can be enabled per interface, but the feature operates with some differences.

At the global level, you can enable BPDU guard on Port Fast-enabled ports by using the **spanning-tree portfast bpduguard default** global configuration command. Spanning tree shuts down ports that are in a Port Fast-operational state. In a valid configuration, Port Fast-enabled ports do not receive BPDUs. Receiving a BPDU on a Port Fast-enabled port signals an invalid configuration, such as the connection of an unauthorized device, and the BPDU guard feature puts the port in the error-disabled state.

At the interface level, you can enable BPDU guard on any port by using the **spanning-tree bpduguard enable** interface configuration command without also enabling the Port Fast feature. When the port receives a BPDU, it is put in the error-disabled state.

The BPDU guard feature provides a secure response to invalid configurations because you must manually put the port back in service. Use the BPDU guard feature in a service-provider network to prevent an access port from participating in the spanning tree.

If your switch is running PVST+, rapid PVST+, or MSTP, you can enable the BPDU guard feature for the entire switch or for an interface.

Understanding BPDU Filtering

The BPDU filtering feature can be globally enabled on the switch or can be enabled per interface, but the feature operates with some differences.

At the global level, you can enable BPDU filtering on Port Fast-enabled ports by using the **spanning-tree portfast bpdupfilter default** global configuration command. This command prevents ports that are in a Port Fast-operational state from sending or receiving BPDUs. The ports still send a few BPDUs at link-up before the switch begins

to filter outbound BPDUs. You should globally enable BPDU filtering on a switch so that hosts connected to these ports do not receive BPDUs. If a BPDU is received on a Port Fast-enabled port, the port loses its Port Fast-operational status, and BPDU filtering is disabled.

At the interface level, you can enable BPDU filtering on any port without also enabling the Port Fast feature by using the **spanning-tree bpdudfilter enable** interface configuration command. This command prevents the port from sending or receiving BPDUs.


Caution: Enabling BPDU filtering on an interface is the same as disabling spanning tree on it and can result in spanning-tree loops.

If your switch is running PVST+, rapid PVST+, or MSTP, you can enable the BPDU filtering feature for the entire switch or for an interface.

Understanding UplinkFast

Switches in hierarchical networks can be grouped into backbone switches, distribution switches, and access switches. Figure 42 shows a complex network where distribution switches and access switches each have at least one redundant link that spanning tree blocks to prevent loops.

Figure 42. Switches in a Hierarchical Network



If a switch loses connectivity, it begins using the alternate paths as soon as the spanning tree selects a new root port. By enabling UplinkFast with the **spanning-tree uplinkfast** global configuration command, you can accelerate the choice of a new root port when a link or switch fails or when the spanning tree reconfigures itself. The root port transitions to the forwarding state immediately without going through the listening and learning states, as it would with the normal spanning-tree procedures. The UplinkFast feature is supported only when the switch is running PVST+. It is not supported when the switch is running rapid PVST+ or MSTP because these protocols use fast convergence and take precedence over UplinkFast.


When the spanning tree reconfigures the new root port, other interfaces flood the network with multicast packets, one for each address that was learned on the interface. You can limit these bursts of multicast traffic by reducing the max-update-rate parameter (the default for this parameter is 150 packets per second). However, if you enter zero, station-learning frames are not generated, so the spanning-tree topology converges more slowly after a loss of connectivity.

Note: UplinkFast is most useful in wiring-closet switches at the access or edge of the network. It is not appropriate for backbone devices. This feature might not be useful for other types of applications.

UplinkFast provides fast convergence after a direct link failure and achieves load balancing between redundant Layer 2 links using uplink groups. An uplink group is a set of Layer 2 interfaces (per VLAN), only one of which is forwarding at any given time. Specifically, an uplink group consists of the root port (which is forwarding) and a set of blocked ports, except for self-looping ports. The uplink group provides an alternate path in case the currently forwarding link fails.


Figure 43 shows an example topology with no link failures. Switch A, the root switch, is connected directly to Switch B over link L1 and to Switch C over link L2. The Layer 2 interface on Switch C that is connected directly to Switch B is in a blocking state.

Figure 43. UplinkFast Example Before Direct Link Failure



If Switch C detects a link failure on the currently active link L2 on the root port (a *direct* link failure), UplinkFast unblocks the blocked port on Switch C and transitions it to the forwarding state without going through the listening and learning states, as shown in Figure 44. This change takes approximately 1 to 5 seconds.

Figure 44. UplinkFast Example After Direct Link Failure



Understanding BackboneFast

BackboneFast detects indirect failures in the core of the backbone. BackboneFast is a complementary technology to the UplinkFast feature, which responds to failures on links directly connected to access switches. BackboneFast optimizes the maximum-age timer, which determines the amount of time the switch stores protocol information received on an interface. When a switch receives an inferior BPDU from the designated port of another switch, the BPDU is a signal that the other switch might have lost its path to the root, and BackboneFast tries to find an alternate path to the root. The BackboneFast feature is supported only when the switch is running PVST+. It is not supported when the switch is running rapid PVST+ or MSTP.

BackboneFast, which is enabled by using the **spanning-tree backbonefast** global configuration command, starts when a root port or blocked port on a switch receives inferior BPDUs from its designated switch. An inferior BPDU identifies a switch that declares itself as both the root bridge and the designated switch. When a switch receives an inferior BPDU, it means that a link to which the switch is not directly connected (an *indirect* link) has failed (that is, the designated bridge has lost its connection to the root switch). Under spanning-tree rules, the switch ignores inferior BPDUs for the configured maximum aging time specified by the **spanning-tree vlan *vlan-id* max-age** global configuration command.


The switch tries to determine if it has an alternate path to the root switch. If the inferior BPDU arrives on a blocked port, the root port and other blocked ports on the switch become alternate paths to the root switch. (Self-looped ports are not considered alternate paths to the root switch.) If the inferior BPDU arrives on the root port, all blocked ports become alternate paths to the root switch. If the inferior BPDU arrives on the root port and there are no blocked ports, the switch assumes that it has lost connectivity to the root switch, causes the maximum aging time on the root port to expire, and becomes the root switch according to normal spanning-tree rules.

If the switch has alternate paths to the root switch, it uses these alternate paths to send a root link query (RLQ) request. The switch sends the RLQ request on all alternate paths to the root switch and waits for an RLQ reply from other switches in the network.

If the switch determines that it still has an alternate path to the root, it expires the maximum aging time on the port that received the inferior BPDU. If all the alternate paths to the root switch indicate that the switch has lost connectivity to the root switch, the switch expires the maximum aging time on the port that received the RLQ reply. If one or more alternate paths can still connect to the root switch, the switch makes all ports on which it received an inferior BPDU its designated ports and moves them from the blocking state (if they were in the blocking state), through the listening and learning states, and into the forwarding state.


Figure 45 shows an example topology with no link failures. Switch A, the root switch, connects directly to Switch B over link L1 and to Switch C over link L2. The Layer 2 interface on Switch C that connects directly to Switch B is in the blocking state.

Figure 45. BackboneFast Example Before Indirect Link Failure




If link L1 fails as shown in Figure 46, Switch C cannot detect this failure because it is not connected directly to link L1. However, because Switch B is directly connected to the root switch over L1, it detects the failure, elects itself the root, and begins sending BPDUs to Switch C, identifying itself as the root. When Switch C receives the inferior BPDUs from Switch B, Switch C assumes that an indirect failure has occurred. At that point, BackboneFast allows the blocked port on Switch C to move immediately to the listening state without waiting for the maximum aging time for the port to expire. BackboneFast then transitions the Layer 2 interface on Switch C to the forwarding state, providing a path from Switch B to Switch A. This switchover takes approximately 30 seconds, twice the Forward Delay time if the default Forward Delay time of 15 seconds is set. Figure 46 shows how BackboneFast reconfigures the topology to account for the failure of link L1.

Figure 46. BackboneFast Example After Indirect Link Failure



If a new switch is introduced into a shared-medium topology as shown in Figure 47, BackboneFast is not activated because the inferior BPDUs did not come from the recognized designated bridge (Switch B). The new switch begins sending inferior BPDUs that indicate it is the root switch. However, the other switches ignore these inferior BPDUs, and the new switch learns that Switch B is the designated bridge to Switch A, the root switch.

Figure 47. Adding a Switch in a Shared-Medium Topology



Understanding EtherChannel Guard

You can use EtherChannel guard to detect an EtherChannel misconfiguration between the switch and a connected device. A misconfiguration can occur if the switch interfaces are configured in an EtherChannel, but the interfaces on the other device are not. A misconfiguration can also occur if the channel parameters are not the same at both ends of the EtherChannel. For EtherChannel configuration guidelines, see the “EtherChannel Configuration Guidelines” section on page 449.

If the switch detects a misconfiguration on the other device, EtherChannel guard places the switch interfaces in the error-disabled state, and this error message appears:

```
PM-4-ERR_DISABLE: Channel-misconfig error detected on [chars], putting [chars] in err-disable state.
```

If your switch is running PVST+, rapid PVST+, or MSTP, you can enable this feature by using the **spanning-tree etherchannel guard misconfig** global configuration command.

Understanding Root Guard

The Layer 2 network of a service provider (SP) can include many connections to switches that are not owned by the SP. In such a topology, the spanning tree can reconfigure itself and select a *customer switch* as the root switch, as shown in Figure 48. You can avoid this situation by enabling root guard on SP switch interfaces that connect to switches in your customer’s network. If spanning-tree calculations cause an interface in the customer network to be selected as the root port, root guard then places the interface in the root-inconsistent (blocked) state to prevent the customer’s switch from becoming the root switch or being in the path to the root.

If a switch outside the SP network becomes the root switch, the interface is blocked (root-inconsistent state), and spanning tree selects a new root switch. The customer’s switch does not become the root switch and is not in the path to the root.

If the switch is operating in multiple spanning-tree (MST) mode, root guard forces the port to be a designated port. If a boundary port is blocked in an internal spanning-tree


(IST) instance because of root guard, the port also is blocked in all MST instances. A boundary port is a port that connects to a LAN, the designated switch of which is either an 802.1D switch or a switch with a different MST region configuration.

Root guard enabled on an interface applies to all the VLANs to which the interface belongs. VLANs can be grouped and mapped to an MST instance.

If your switch is running PVST+, rapid PVST+, or MSTP, you can enable this feature by using the **spanning-tree guard root** interface configuration command.

Caution: Misuse of the root-guard feature can cause a loss of connectivity.

Figure 48. Root Guard in a Service-Provider Network



Understanding Loop Guard

You can use loop guard to prevent alternate or root ports from becoming designated ports because of a failure that leads to a unidirectional link. This feature is most effective when it is configured on the entire switched network. Loop guard prevents alternate and root ports from becoming designated ports, and spanning tree does not send BPDUs on root or alternate ports.

If your switch is running PVST+, rapid PVST+, or MSTP, you can enable this feature by using the **spanning-tree loopguard default** global configuration command.

When the switch is operating in PVST+ or rapid-PVST+ mode, loop guard prevents alternate and root ports from becoming designated ports, and spanning tree does not send BPDUs on root or alternate ports.

When the switch is operating in MST mode, BPDUs are not sent on nonboundary ports only if the port is blocked by loop guard in all MST instances. On a boundary port, loop guard blocks the port in all MST instances.

Configuring Optional Spanning-Tree Features

These sections describe how to configure optional spanning-tree features:

- Default Optional Spanning-Tree Configuration, on page 235
- Optional Spanning-Tree Configuration Guidelines, on page 235

- Enabling Port Fast, on page 235 (optional)
- Enabling BPDU Guard, on page 236 (optional)
- Enabling BPDU Filtering, on page 237 (optional)
- Enabling UplinkFast for Use with Redundant Links, on page 238 (optional)
- Enabling BackboneFast, on page 239 (optional)
- Enabling EtherChannel Guard, on page 239 (optional)
- Enabling Root Guard, on page 240 (optional)
- Enabling Loop Guard, on page 240 (optional)

Default Optional Spanning-Tree Configuration

Table 45 shows the default optional spanning-tree configuration.

Table 45. Default Optional Spanning-Tree Configuration

Feature	Default Setting
Port Fast, BPDU filtering, BPDU guard	Enabled on internal ports 1-14. Disabled on ports 15-20 (unless they are individually configured per interface).
UplinkFast	Globally disabled.
CSUF	Disabled on all interfaces.
BackboneFast	Globally disabled.
EtherChannel guard	Globally enabled.
Root guard	Disabled on all interfaces.
Loop guard	Disabled on all interfaces.

Optional Spanning-Tree Configuration Guidelines

The UplinkFast and BackboneFast, features are not supported with the rapid PVST+ or the MSTP.

Enabling Port Fast

A port with the Port Fast feature enabled is moved directly to the spanning-tree forwarding state without waiting for the standard forward-time delay.

Caution: Use Port Fast *only* when connecting a single end station to an access or trunk port. Enabling this feature on a port connected to a switch or hub could prevent spanning tree from detecting and disabling loops in your network, which could cause broadcast storms and address-learning problems.

If you enable the voice VLAN feature, the Port Fast feature is automatically enabled. When you disable voice VLAN, the Port Fast feature is not automatically disabled.

You can enable this feature if your switch is running PVST+, rapid PVST+, or MSTP.

Beginning in privileged EXEC mode, follow these steps to enable Port Fast. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify an interface to configure.
3.	spanning-tree portfast [trunk]	<p>Enable Port Fast on an access port connected to a single workstation or server. By specifying the trunk keyword, you can enable Port Fast on a trunk port.</p> <p>Caution: Make sure that there are no loops in the network between the trunk port and the workstation or server before you enable Port Fast on a trunk port.</p> <p>By default, Port Fast is enabled on ports 1-14 and disabled on ports 15-20.</p>
4.	end	Return to privileged EXEC mode.
5.	show spanning-tree interface <i>interface-id</i> portfast	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Note: You can use the **spanning-tree portfast default** global configuration command to globally enable the Port Fast feature on all nontrunking ports.

To disable the Port Fast feature, use the **spanning-tree portfast disable** interface configuration command.

Enabling BPDU Guard

When you globally enable BPDU guard on ports that are Port Fast-enabled (the ports are in a Port Fast-operational state), spanning tree shuts down Port Fast-enabled ports that receive BPDUs.

In a valid configuration, Port Fast-enabled ports do not receive BPDUs. Receiving a BPDU on a Port Fast-enabled port signals an invalid configuration, such as the connection of an unauthorized device, and the BPDU guard feature puts the port in the error-disabled state. The BPDU guard feature provides a secure response to invalid configurations because you must manually put the port back in service. Use the BPDU guard feature in a service-provider network to prevent an access port from participating in the spanning tree.

Caution: Configure Port Fast only on ports that connect to end stations; otherwise, an accidental topology loop could cause a data packet loop and disrupt switch and network operation.

You can also use the **spanning-tree bpduguard enable** interface configuration command to enable BPDU guard on any port without also enabling the Port Fast feature. When the port receives a BPDU, it is put in the error-disabled state.

You can enable the BPDU guard feature if your switch is running PVST+, rapid PVST+, or MSTP.

Beginning in privileged EXEC mode, follow these steps to globally enable the BPDU guard feature. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree portfast bpduguard default	Globally enable BPDU guard. By default, BPDU guard is disabled.
3.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface connected to an end station.
4.	spanning-tree portfast	Enable the Port Fast feature.
5.	end	Return to privileged EXEC mode.
6.	show running-config	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable BPDU guard, use the **no spanning-tree portfast bpduguard default** global configuration command.

You can override the setting of the **no spanning-tree portfast bpduguard default** global configuration command by using the **spanning-tree bpduguard enable** interface configuration command.

Enabling BPDU Filtering

When you globally enable BPDU filtering on Port Fast-enabled ports, it prevents ports that are in a Port Fast-operational state from sending or receiving BPDUs. The ports still send a few BPDUs at link-up before the switch begins to filter outbound BPDUs. You should globally enable BPDU filtering on a switch so that hosts connected to these ports do not receive BPDUs. If a BPDU is received on a Port Fast-enabled port, the port loses its Port Fast-operational status, and BPDU filtering is disabled.

Note: BPDU filtering is enabled by default on the blade ports (ports 1-14). It is not globally enabled. You may want to change this configuration to BPDU guard to ensure that the internal blade ports do not start to bridge ethernet traffic and participate in STP topology.

Caution: Configure Port Fast only on ports that connect to end stations; otherwise, an accidental topology loop could cause a data packet loop and disrupt switch and network operation.

You can also use the **spanning-tree bpdudfilter enable** interface configuration command to enable BPDU filtering on any port without also enabling the Port Fast feature. This command prevents the port from sending or receiving BPDUs.

Caution: Enabling BPDU filtering on an interface is the same as disabling spanning tree on it and can result in spanning-tree loops.

You can enable the BPDU filtering feature if your switch is running PVST+, rapid PVST+, or MSTP.

Beginning in privileged EXEC mode, follow these steps to globally enable the BPDU filtering feature. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree portfast bpdudfilter default	Globally enable BPDU filtering. By default, BPDU filtering is disabled.

Step	Command	Purpose
3.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface connected to an end station.
4.	spanning-tree portfast	Enable the Port Fast feature.
5.	end	Return to privileged EXEC mode.
6.	show running-config	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable BPDU filtering, use the **no spanning-tree portfast bpdupfilter default** global configuration command.

You can override the setting of the **no spanning-tree portfast bpdupfilter default** global configuration command by using the **spanning-tree bpdupfilter enable** interface configuration command.

Enabling UplinkFast for Use with Redundant Links

UplinkFast cannot be enabled on VLANs that have been configured for switch priority. To enable UplinkFast on a VLAN with switch priority configured, first restore the switch priority on the VLAN to the default value by using the **no spanning-tree vlan *vlan-id* priority** global configuration command.

Note: When you enable UplinkFast, it affects all VLANs on the switch. You cannot configure UplinkFast on an individual VLAN.

The UplinkFast feature is supported only when the switch is running PVST+. It is not supported when the switch is running rapid PVST+ or MSTP.

Beginning in privileged EXEC mode, follow these steps to enable UplinkFast. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree uplinkfast [max-update-rate <i>pkts-per-second</i>]	Enable UplinkFast. (Optional) For <i>pkts-per-second</i> , the range is 0 to 32000 packets per second; the default is 150. If you set the rate to 0, station-learning frames are not generated, and the spanning-tree topology converges more slowly after a loss of connectivity.
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree summary	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

When UplinkFast is enabled, the switch priority of all VLANs is set to 49152. If you change the path cost to a value less than 3000 and you enable UplinkFast or UplinkFast is already enabled, the path cost of all interfaces and VLAN trunks is increased by 3000 (if you change the path cost to 3000 or above, the path cost is not altered). The changes to the switch priority and the path cost reduces the chance that the switch will become the root switch.

When UplinkFast is disabled, the switch priorities of all VLANs and path costs of all interfaces are set to default values if you did not modify them from their defaults.

To return the update packet rate to the default setting, use the **no spanning-tree uplinkfast max-update-rate** global configuration command. To disable UplinkFast, use the **no spanning-tree uplinkfast** command.

Enabling BackboneFast

You can enable BackboneFast to detect indirect link failures and to start the spanning-tree reconfiguration sooner.

Note: If you use BackboneFast, you must enable it on all switches in the network. BackboneFast is not supported on Token Ring VLANs. This feature is supported for use with third-party switches.

The BackboneFast feature is supported only when the switch is running PVST+. It is not supported when the switch is running rapid PVST+ or MSTP.

Beginning in privileged EXEC mode, follow these steps to enable BackboneFast. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree backbonefast	Enable BackboneFast.
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree summary	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the BackboneFast feature, use the **no spanning-tree backbonefast** global configuration command.

Enabling EtherChannel Guard

You can enable EtherChannel guard to detect an EtherChannel misconfiguration that causes a loop.

You can enable this feature if your switch is running PVST+, rapid PVST+, or MSTP.

Beginning in privileged EXEC mode, follow these steps to enable EtherChannel guard. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	spanning-tree etherchannel guard misconfig	Enable EtherChannel guard.
3.	end	Return to privileged EXEC mode.
4.	show spanning-tree summary	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the EtherChannel guard feature, use the **no spanning-tree etherchannel guard misconfig** global configuration command.

You can use the **show interfaces status err-disabled** privileged EXEC command to determine which switch ports are disabled because of an EtherChannel misconfiguration. On the remote device, you can enter the **show etherchannel summary** privileged EXEC command to verify the EtherChannel configuration.

After the configuration is corrected, enter the **shutdown** and **no shutdown** interface configuration commands on the port-channel interfaces that were misconfigured.

Enabling Root Guard

Root guard enabled on an interface applies to all the VLANs to which the interface belongs.

Do not enable the root guard on interfaces to be used by the UplinkFast feature. With UplinkFast, the backup interfaces (in the blocked state) replace the root port in the case of a failure. However, if root guard is also enabled, all the backup interfaces used by the UplinkFast feature are placed in the root-inconsistent state (blocked) and are prevented from reaching the forwarding state.

Note: You cannot enable both root guard and loop guard at the same time.

You can enable this feature if your switch is running PVST+, rapid PVST+, or MSTP.

Beginning in privileged EXEC mode, follow these steps to enable root guard on an interface. This procedure is optional.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify an interface to configure.
3.	spanning-tree guard root	Enable root guard on the interface. By default, root guard is disabled on all interfaces.
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable root guard, use the **no spanning-tree guard** interface configuration command.

Enabling Loop Guard

You can use loop guard to prevent alternate or root ports from becoming designated ports because of a failure that leads to a unidirectional link. This feature is most effective when it is configured on the entire switched network. Loop guard operates only on ports that are considered point-to-point by the spanning tree.

Note: You cannot enable both loop guard and root guard at the same time.

You can enable this feature if your switch is running PVST+, rapid PVST+, or MSTP.

Beginning in privileged EXEC mode, follow these steps to enable loop guard. This procedure is optional.

Step	Command	Purpose
1.	show spanning-tree active or show spanning-tree mst	Determine which ports are alternate or root ports.
2.	configure terminal	Enter global configuration mode.

Step	Command	Purpose
3.	spanning-tree loopguard default	Enable loop guard. By default, loop guard is disabled.
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To globally disable loop guard, use the **no spanning-tree loopguard default** global configuration command. You can override the setting of the **no spanning-tree loopguard default** global configuration command by using the **spanning-tree guard loop** interface configuration command.

Displaying the Spanning-Tree Status

To display the spanning-tree status, use one or more of the privileged EXEC commands in Table 46:

Table 46. Commands for Displaying the Spanning-Tree Status

Command	Purpose
show spanning-tree active	Displays spanning-tree information on active interfaces only.
show spanning-tree detail	Displays a detailed summary of interface information.
show spanning-tree interface <i>interface-id</i>	Displays spanning-tree information for the specified interface.
show spanning-tree mst interface <i>interface-id</i>	Displays MST information for the specified interface.
show spanning-tree summary [totals]	Displays a summary of port states or displays the total lines of the spanning-tree state section.

You can clear spanning-tree counters by using the **clear spanning-tree [interface *interface-id*]** privileged EXEC command.

For information about other keywords for the **show spanning-tree** privileged EXEC command, refer to the *command reference* for this release.

Chapter 13. Configuring VLANs

This chapter describes how to configure normal-range VLANs (VLAN IDs 1 to 1005) and extended-range VLANs (VLAN IDs 1006 to 4094) on your switch. It includes information about VLAN modes and the VLAN Membership Policy Server (VMPS).

Note: For complete syntax and usage information for the commands used in this chapter, refer to the *command reference* for this release.

The chapter includes these sections:

- Understanding VLANs, on page 243
- Configuring Normal-Range VLANs, on page 245
- Configuring Extended-Range VLANs, on page 252
- Displaying VLANs, on page 254
- Configuring VLAN Trunks, on page 255
- Configuring VMPS, on page 266


Understanding VLANs

A VLAN is a switched network that is logically segmented by function, project team, or application, without regard to the physical locations of the users. VLANs have the same attributes as physical LANs, but you can group end stations even if they are not physically located on the same LAN segment. Any switch port can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to end stations in the VLAN. Each VLAN is considered a logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a router or bridge as shown in Figure 49. Because a VLAN is considered a separate logical network, it contains its own bridge Management Information Base (MIB) information and can support its own implementation of spanning tree. See Chapter 10 “Configuring STP” and Chapter 11 “Configuring MSTP.”

Note: Before you create VLANs, you must decide whether to use VLAN Trunking Protocol (VTP) to maintain global VLAN configuration for your network. For more information on VTP, see Chapter 14 “Configuring VTP.”

Figure 49 shows an example of VLANs segmented into logically defined networks.

Figure 49. VLANs as Logically Defined Networks



VLANs are often associated with IP subnetworks. For example, all the end stations in a particular IP subnet belong to the same VLAN. Interface VLAN membership on the switch is assigned manually on an interface-by-interface basis. When you assign switch interfaces to VLANs by using this method, it is known as interface-based, or static, VLAN membership.

Supported VLANs

The switch supports 250 VLANs. Refer to the release notes for the list of switches that support each image. VLANs are identified with a number from 1 to 4094. VLAN IDs 1002 through 1005 are reserved for Token Ring and FDDI VLANs. VTP only learns normal-range VLANs, with VLAN IDs 1 to 1005; VLAN IDs greater than 1005 are extended-range VLANs and are not stored in the VLAN database. The switch must be in VTP transparent mode when you create VLAN IDs from 1006 to 4094.

The switch supports per-VLAN spanning-tree plus (PVST+) and rapid PVST+ with a maximum of 64 spanning-tree instances. One spanning-tree instance is allowed per VLAN. See the “Normal-Range VLAN Configuration Guidelines” section on page 247 for more information about the number of spanning-tree instances and the number of VLANs. The switch supports IEEE 802.1Q trunking for sending VLAN traffic over Ethernet ports.

VLAN Port Membership Modes

You configure a port to belong to a VLAN by assigning a membership mode that determines the kind of traffic the port carries and the number of VLANs to which it can belong. Table 47 lists the membership modes and membership and VTP characteristics.

Table 47. Port Membership Modes

Membership Mode	VLAN Membership Characteristics	VTP Characteristics
Static-access	A static-access port can belong to one VLAN and is manually assigned to that VLAN. For more information, see the “Assigning Static-Access Ports to a VLAN” section on page 251.	VTP is not required. If you do not want VTP to globally propagate information, set the VTP mode to transparent to disable VTP. To participate in VTP, there must be at least one trunk port on the switch connected to a trunk port of a second switch.
802.1Q trunk	A trunk port is a member of all VLANs by default, including extended-range VLANs, but membership can be limited by configuring the allowed-VLAN list. You can also modify the pruning-eligible list to block flooded traffic to VLANs on trunk ports that are included in the list. For information about configuring trunk ports, see the “Configuring an Ethernet Interface as a Trunk Port” section on page 258.	VTP is recommended but not required. VTP maintains VLAN configuration consistency by managing the addition, deletion, and renaming of VLANs on a network-wide basis. VTP exchanges VLAN configuration messages with other switches over trunk links.
Dynamic access	<p>A dynamic-access port can belong to one normal-range VLAN (VLAN ID 1 to 1005) and is dynamically assigned by a VMPS. The VMPS can be a Catalyst 5000 or Catalyst 6000 series switch, for example, but never a Cisco Systems Intelligent Gigabit Ethernet Switch Module.</p> <p>You can have dynamic-access ports and trunk ports on the same switch, but you must connect the dynamic-access port to an end station and not to another switch.</p> <p>For configuration information, see the “Configuring Dynamic Access Ports on VMPS Clients” section on page 270.</p>	<p>VTP is required.</p> <p>Configure the VMPS and the client with the same VTP domain name.</p> <p>You can change the reconfirmation interval and retry count on the VMPS client switch.</p>

For more detailed definitions of the modes and their functions, see Table 50 on page 256.

When a port belongs to a VLAN, the switch learns and manages the addresses associated with the port on a per-VLAN basis. For more information, see the “Managing the MAC Address Table” section on page 110.

Configuring Normal-Range VLANs

Normal-range VLANs are VLANs with VLAN IDs 1 to 1005. If the switch is in VTP server or transparent mode, you can add, modify or remove configurations for VLANs 2 to 1001 in the VLAN database. (VLAN IDs 1 and 1002 to 1005 are automatically created and cannot be removed.)

Note: When the switch is in VTP transparent mode, you can also create extended-range VLANs (VLANs with IDs from 1006 to 4094), but these VLANs are not saved in the VLAN database. See the “Configuring Extended-Range VLANs” section on page 252.

Configurations for VLAN IDs 1 to 1005 are written to the file *vlan.dat* (VLAN database), and you can display them by entering the **show vlan** privileged EXEC command. The *vlan.dat* file is stored in nonvolatile RAM (NVRAM).

Caution: You can cause inconsistency in the VLAN database if you attempt to manually delete the *vlan.dat* file. If you want to modify the VLAN configuration, use the commands described in these sections and in the *command reference* for this release. To change the VTP configuration, see Chapter 14 “Configuring VTP.”

You use the interface configuration mode to define the port membership mode and to add and remove ports from VLANs. The results of these commands are written to the running-configuration file, and you can display the file by entering the **show running-config** privileged EXEC command.

You can set these parameters when you create a new normal-range VLAN or modify an existing VLAN in the VLAN database:

- VLAN ID
- VLAN name
- VLAN type (Ethernet, Fiber Distributed Data Interface [FDDI], FDDI network entity title [NET], TrBRF, or TrCRF, Token Ring, Token Ring-Net)
- VLAN state (active or suspended)
- Maximum transmission unit (MTU) for the VLAN
- Security Association Identifier (SAID)
- Bridge identification number for TrBRF VLANs
- Ring number for FDDI and TrCRF VLANs
- Parent VLAN number for TrCRF VLANs
- Spanning Tree Protocol (STP) type for TrCRF VLANs
- VLAN number to use when translating from one VLAN type to another

Note: This section does not provide configuration details for most of these parameters. For complete information on the commands and parameters that control VLAN configuration, refer to the *command reference* for this release.

This section includes information about these topics about normal-range VLANs:

- Token Ring VLANs, on page 246
- Normal-Range VLAN Configuration Guidelines, on page 247
- VLAN Configuration Mode Options, on page 247
- Saving VLAN Configuration, on page 248
- Default Ethernet VLAN Configuration, on page 249
- Creating or Modifying an Ethernet VLAN, on page 249
- Deleting a VLAN, on page 251
- Assigning Static-Access Ports to a VLAN, on page 251

Token Ring VLANs

Although the switch does not support Token Ring connections, a remote device such as a Catalyst 5000 series switch with Token Ring connections could be managed from one of the supported switches. Switches running VTP version 2 advertise information about these Token Ring VLANs:

- Token Ring TrBRF VLANs

- Token Ring TrCRF VLANs

For more information on configuring Token Ring VLANs, refer to the *Catalyst 5000 Series Software Configuration Guide*.

Normal-Range VLAN Configuration Guidelines

Follow these guidelines when creating and modifying normal-range VLANs in your network:

- See Table 47 for the maximum number of supported VLANs per switch model. On a switch supporting 250 VLANs, if VTP reports that there are 250 active VLANs, four of the active VLANs (1002 to 1005) are reserved for Token Ring and FDDI.
- Normal-range VLANs are identified with a number between 1 and 1001. VLAN numbers 1002 through 1005 are reserved for Token Ring and FDDI VLANs.
- VLAN configuration for VLANs 1 to 1005 are always saved in the VLAN database. If VTP mode is transparent, VTP and VLAN configuration is also saved in the switch running configuration file.
- The switch also supports VLAN IDs 1006 through 4094 in VTP transparent mode (VTP disabled). These are extended-range VLANs and configuration options are limited. Extended-range VLANs are not saved in the VLAN database. See the “Configuring Extended-Range VLANs” section on page 252.
- Before you can create a VLAN, the switch must be in VTP server mode or VTP transparent mode. If the switch is a VTP server, you must define a VTP domain or VTP will not function.
- The switch does not support Token Ring or FDDI media. The switch does not forward FDDI, FDDI-Net, TrCRF, or TrBRF traffic, but it does propagate the VLAN configuration through VTP.
- The switch supports 64 spanning-tree instances. If a switch has more active VLANs than supported spanning-tree instances, spanning tree can be enabled on 64 VLANs and is disabled on the remaining VLANs. If you have already used all available spanning-tree instances on a switch, adding another VLAN anywhere in the VTP domain creates a VLAN on that switch that is not running spanning-tree. If you have the default allowed list on the trunk ports of that switch (which is to allow all VLANs), the new VLAN is carried on all trunk ports. Depending on the topology of the network, this could create a loop in the new VLAN that would not be broken, particularly if there are several adjacent switches that all have run out of spanning-tree instances. You can prevent this possibility by setting allowed lists on the trunk ports of switches that have used up their allocation of spanning-tree instances.

If the number of VLANs on the switch exceeds the number of supported spanning tree instances, we recommend that you configure the IEEE 802.1s Multiple STP (MSTP) on your switch to map multiple VLANs to a single STP instance. For more information about MSTP, see Chapter 11 “Configuring MSTP.”

VLAN Configuration Mode Options

You can configure normal-range VLANs (with VLAN IDs 1 to 1005) by using these two configuration modes:

- VLAN Configuration in config-vlan Mode, on page 248
You access config-vlan mode by entering the **vlan *vlan-id*** global configuration command.
- VLAN Configuration in VLAN Configuration Mode, on page 248

You access VLAN database configuration mode by entering the **vlan database** privileged EXEC command.

VLAN Configuration in config-vlan Mode

To access config-vlan mode, enter the **vlan** global configuration command with a VLAN ID. Enter a new VLAN ID to create a VLAN or with an existing VLAN ID to modify the VLAN. You can use the default VLAN configuration (Table 48) or enter multiple commands to configure the VLAN. For more information about commands available in this mode, refer to the **vlan** global configuration command description in the *command reference* for this release. When you have finished the configuration, you must exit config-vlan mode for the configuration to take effect. To display the VLAN configuration, enter the **show vlan** privileged EXEC command.

You must use this config-vlan mode when creating extended-range VLANs (VLAN IDs greater than 1005). See the “Configuring Extended-Range VLANs” section on page 252.

VLAN Configuration in VLAN Configuration Mode

To access VLAN configuration mode, enter the **vlan database** privileged EXEC command. Then enter the **vlan** command with a new VLAN ID to create a VLAN or with an existing VLAN ID to modify the VLAN. You can use the default VLAN configuration (Table 48) or enter multiple commands to configure the VLAN. For more information about keywords available in this mode, refer to the **vlan** VLAN configuration command description in the *command reference* for this release. When you have finished the configuration, you must enter **apply** or **exit** for the configuration to take effect. When you enter the **exit** command, it applies all commands and updates the VLAN database. VTP messages are sent to other switches in the VTP domain, and the privileged EXEC mode prompt appears.

Saving VLAN Configuration

The configurations of VLAN IDs 1 to 1005 are always saved in the VLAN database (vlan.dat file). If VTP mode is transparent, they are also saved in the switch running configuration file and you can enter the **copy running-config startup-config** privileged EXEC command to save the configuration in the startup configuration file. You can use the **show running-config vlan** privileged EXEC command to display the switch running configuration file. To display the VLAN configuration, enter the **show vlan** privileged EXEC command.

When you save VLAN and VTP information (including extended-range VLAN configuration information) in the startup configuration file and reboot the switch, the switch configuration is determined as follows:

- If the VTP mode is transparent in the startup configuration, and the VLAN database and the VTP domain name from the VLAN database matches that in the startup configuration file, the VLAN database is ignored (cleared), and the VTP and VLAN configurations in the startup configuration file are used. The VLAN database revision number remains unchanged in the VLAN database.
- If the VTP mode or domain name in the startup configuration does not match the VLAN database, the domain name and VTP mode and configuration for the first 1005 VLANs use the VLAN database information.
- If VTP mode is server, the domain name and VLAN configuration for the first 1005 VLANs use the VLAN database information.

Caution: If the VLAN database configuration is used at startup and the startup configuration file contains extended-range VLAN configuration, this information is lost when the system boots up.

Default Ethernet VLAN Configuration

Table 48 shows the default configuration for Ethernet VLANs.

Note: The switch supports Ethernet interfaces exclusively. Because FDDI and Token Ring VLANs are not locally supported, you only configure FDDI and Token Ring media-specific characteristics for VTP global advertisements to other switches.

Table 48. Ethernet VLAN Defaults and Ranges

Parameter	Default	Range
VLAN ID	1 (for the internal 100 Mbps management module ports)	No range
	2 (for the internal 1000 Mbps ports and the external 10/100/1000 Mbps ports)	No range
VLAN name	For VLAN 1: <i>default</i> For VLAN 2: <i>operational</i>	No range
802.10 SAID	100001 (100000 plus the VLAN ID)	1–4294967294
MTU size	1500	1500–18190
Translational bridge 1	0	0–1005
Translational bridge 2	0	0–1005
VLAN state	For VLAN 1: active For VLAN 2: active	active, suspend
Remote SPAN	disabled	enabled, disabled

Creating or Modifying an Ethernet VLAN

Each Ethernet VLAN in the VLAN database has a unique, 4-digit ID that can be a number from 1 to 1001. VLAN IDs 1002 to 1005 are reserved for Token Ring and FDDI VLANs. To create a normal-range VLAN to be added to the VLAN database, assign a number and name to the VLAN.

Note: When the switch is in VTP transparent mode, you can assign VLAN IDs greater than 1006, but they are not added to the VLAN database. See the “Configuring Extended-Range VLANs” section on page 252.

For the list of default parameters that are assigned when you add a VLAN, see the “Configuring Normal-Range VLANs” section on page 245.

Beginning in privileged EXEC mode, follow these steps to use config-vlan mode to create or modify an Ethernet VLAN:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	vlan <i>vlan-id</i>	Enter a VLAN ID, and enter config-vlan mode. Enter a new VLAN ID to create a VLAN, or enter an existing VLAN ID to modify a VLAN. Note: The available VLAN ID range for this command is 1 to 4094. For information about adding VLAN IDs greater than 1005 (extended-range VLANs), see the “Configuring Extended-Range VLANs” section on page 252.

Step	Command	Purpose
3.	name <i>vlan-name</i>	(Optional) Enter a name for the VLAN. If no name is entered for the VLAN, the default is to append the <i>vlan-id</i> with leading zeros to the word VLAN. For example, VLAN0004 is a default VLAN name for VLAN 4.
4.	mtu <i>mtu-size</i>	(Optional) Change the MTU size (or other VLAN characteristic).
5.	remote-span	(Optional) Configure the VLAN as the RSPAN VLAN for a remote SPAN session. For more information on remote SPAN, see Chapter 19 "Configuring SPAN and RSPAN."
6.	end	Return to privileged EXEC mode.
7.	show vlan { name <i>vlan-name</i> / id <i>vlan-id</i> }	Verify your entries.
8.	copy running-config startup config	(Optional) If the switch is in VTP transparent mode, the VLAN configuration is saved in the running configuration file as well as in the VLAN database. This saves the configuration in the switch startup configuration file.

To return the VLAN name to the default settings, use the **no vlan name**, **no vlan mtu**, or **no remote span** config-vlan commands.

This example shows how to use config-vlan mode to create Ethernet VLAN 20, name it *test20*, and add it to the VLAN database:

```
Switch# configure terminal
Switch(config)# vlan 20
Switch(config-vlan)# name test20
Switch(config-vlan)# end
```

Beginning in privileged EXEC mode, follow these steps to use VLAN configuration mode to create or modify an Ethernet VLAN:

Step	Command	Purpose
1.	vlan database	Enter VLAN database configuration mode.
2.	vlan <i>vlan-id</i> name <i>vlan-name</i>	Add an Ethernet VLAN by assigning a number to it. The range is 1 to 1001; do not enter leading zeros. If no name is entered for the VLAN, the default is to append the <i>vlan-id</i> with leading zeros to the word VLAN. For example, VLAN0004 is a default VLAN name for VLAN 4.
3.	vlan <i>vlan-id</i> mtu <i>mtu-size</i>	(Optional) To modify a VLAN, identify the VLAN and change a characteristic, such as the MTU size.
4.	exit	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
5.	show vlan { name <i>vlan-name</i> / id <i>vlan-id</i> }	Verify your entries.
6.	copy running-config startup config	(Optional) If the switch is in VTP transparent mode, the VLAN configuration is saved in the running configuration file as well as in the VLAN database. This saves the configuration in the switch startup configuration file.

Note: You cannot configure an RSPAN VLAN in VLAN database configuration mode.

To return the VLAN name to the default settings, use the **no vlan *vlan-id* name** or **no vlan *vlan-id* mtu** VLAN configuration command.

This example shows how to use VLAN database configuration mode to create Ethernet VLAN 20, name it *test20*, and add it to the VLAN database:

```
Switch# vlan database
Switch(vlan)# vlan 20 name test20
Switch(vlan)# exit
APPLY completed.
Exiting....
Switch#
```

Deleting a VLAN

When you delete a VLAN from a switch that is in VTP server mode, the VLAN is removed from the VLAN database for all switches in the VTP domain. When you delete a VLAN from a switch that is in VTP transparent mode, the VLAN is deleted only on that specific switch.

You cannot delete the default VLANs for the different media types: Ethernet VLAN 1 and FDDI or Token Ring VLANs 1002 to 1005.

Caution: When you delete a VLAN, any ports assigned to that VLAN become inactive. They remain associated with the VLAN (and thus inactive) until you assign them to a new VLAN.

Beginning in privileged EXEC mode, follow these steps to delete a VLAN on the switch by using global configuration mode:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	no vlan <i>vlan-id</i>	Remove the VLAN by entering the VLAN ID.
3.	end	Return to privileged EXEC mode.
4.	show vlan brief	Verify the VLAN removal.
5.	copy running-config startup config	(Optional) If the switch is in VTP transparent mode, the VLAN configuration is saved in the running configuration file as well as in the VLAN database. This saves the configuration in the switch startup configuration file.

To delete a VLAN in VLAN database configuration mode, use the **vlan database** privileged EXEC command to enter VLAN database configuration mode and the **no vlan *vlan-id*** VLAN configuration command.

Assigning Static-Access Ports to a VLAN

You can assign a static-access port to a VLAN without having VTP globally propagate VLAN configuration information by disabling VTP (VTP transparent mode). If you are assigning a port on a cluster member switch to a VLAN, first use the **rcommand** privileged EXEC command to log in to the member switch.

Note: If you assign an interface to a VLAN that does not exist, the new VLAN is created. (See the “Creating or Modifying an Ethernet VLAN” section on page 249.)

Beginning in privileged EXEC mode, follow these steps to assign a port to a VLAN in the VLAN database:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode
2.	interface <i>interface-id</i>	Enter the interface to be added to the VLAN.
3.	switchport mode access	Define the VLAN membership mode for the port (Layer 2 access port).
4.	switchport access vlan <i>vlan-id</i>	Assign the port to a VLAN. Valid VLAN IDs are 1 to 4094.
5.	end	Return to privileged EXEC mode.
6.	show running-config interface <i>interface-id</i>	Verify the VLAN membership mode of the interface.
7.	show interfaces <i>interface-id</i> switchport	Verify your entries in the <i>Administrative Mode</i> and the <i>Access Mode VLAN</i> fields of the display.
8.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return an interface to its default configuration, use the **default interface** *interface-id* interface configuration command.

This example shows how to configure Gigabit Ethernet interface 0/17 as an access port in VLAN 2:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan 2
Switch(config-if)# end
Switch#
```

Configuring Extended-Range VLANs

When the switch is in VTP transparent mode (VTP disabled), you can create extended-range VLANs (in the range 1006 to 4094). Extended-range VLANs enable service providers to extend their infrastructure to a greater number of customers. The extended-range VLAN IDs are allowed for any switchport commands that allow VLAN IDs. You always use config-vlan mode (accessed by entering the **vlan** *vlan-id* global configuration command) to configure extended-range VLANs. The extended range is not supported in VLAN database configuration mode (accessed by entering the **vlan database** privileged EXEC command).

Extended-range VLAN configurations are not stored in the VLAN database, but because VTP mode is transparent, they are stored in the switch running configuration file, and you can save the configuration in the startup configuration file by using the **copy running-config startup-config** privileged EXEC command.

Note: Although the switch supports 4094 VLAN IDs, see the “Supported VLANs” section on page 244 for the actual number of VLANs supported.

This section includes this information about extended-range VLANs:

- Default VLAN Configuration, on page 253
- Extended-Range VLAN Configuration Guidelines, on page 253
- Creating an Extended-Range VLAN, on page 253

Default VLAN Configuration

See Table 48 on page 249 for the default configuration for Ethernet VLANs. You can change only the MTU size on extended-range VLANs; all other characteristics must remain at the default state.

Extended-Range VLAN Configuration Guidelines

Follow these guidelines when creating extended-range VLANs:

- To add an extended-range VLAN, you must use the **vlan** *vlan-id* global configuration command and access config-vlan mode. You cannot add extended-range VLANs in VLAN database configuration mode (accessed by entering the **vlan database** privileged EXEC command).
- VLAN IDs in the extended range are not saved in the VLAN database and are not recognized by VTP.
- You cannot include extended-range VLANs in the pruning eligible range.
- The switch must be in VTP transparent mode when you create extended-range VLANs. If VTP mode is server or client, an error message is generated, and the extended-range VLAN is rejected.
- You can set the VTP mode to transparent in global configuration mode or in VLAN database configuration mode. See the “Disabling VTP (VTP Transparent Mode)” section on page 285. You should save this configuration to the startup configuration so that the switch will boot up in VTP transparent mode. Otherwise, you will lose extended-range VLAN configuration if the switch resets.
- VLANs in the extended range are not supported by VQP. They cannot be configured by VMPS.
- STP is enabled by default on extended-range VLANs, but you can disable it by using the **no spanning-tree vlan** *vlan-id* global configuration command. When the maximum number of spanning-tree instances (64) are on the switch, spanning tree is disabled on any newly created VLANs. If the number of VLANs on the switch exceeds the maximum number of spanning tree instances, we recommend that you configure the IEEE 802.1s Multiple STP (MSTP) on your switch to map multiple VLANs to a single STP instance. For more information about MSTP, see Chapter 11 “Configuring MSTP.”

Creating an Extended-Range VLAN

You create an extended-range VLAN in global configuration mode by entering the **vlan** global configuration command with a VLAN ID from 1006 to 4094. This command accesses the config-vlan mode. The extended-range VLAN has the default Ethernet VLAN characteristics (see Table 48) and the MTU size is the only parameter you can change. Refer to the description of the **vlan** global configuration command in the command reference for defaults of all parameters. If you enter an extended-range VLAN ID when the switch is not in VTP transparent mode, an error message is generated when you exit from config-vlan mode, and the extended-range VLAN is not created.

Extended-range VLANs are not saved in the VLAN database; they are saved in the switch running configuration file. You can save the extended-range VLAN configuration in the switch startup configuration file by using the **copy running-config startup-config** privileged EXEC command.

Note: Before you create an extended-range VLAN, you can verify that the VLAN ID is not used internally by entering the **show vlan internal usage** privileged EXEC command. If the VLAN ID is used internally and you want to free it up, go to

the “Displaying VLANs” section on page 254 before creating the extended-range VLAN.

Beginning in privileged EXEC mode, follow these steps to create an extended-range VLAN:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	vtp mode transparent	Configure the switch for VTP transparent mode, disabling VTP.
3.	vlan <i>vlan-id</i>	Enter an extended-range VLAN ID and enter config-vlan mode. The range is 1006 to 4094.
4.	mtu <i>mtu-size</i>	(Optional) Modify the VLAN by changing the MTU size. Note: Although all commands appear in the CLI help in config-vlan mode, only the mtu <i>mtu-size</i> command is supported for extended-range VLANs.
5.	end	Return to privileged EXEC mode.
6.	show vlan id <i>vlan-id</i>	Verify that the VLAN has been created.
7.	copy running-config startup config	Save your entries in the switch startup configuration file. To save extended-range VLAN configurations, you need to save the VTP transparent mode configuration and the extended-range VLAN configuration in the switch startup configuration file. Otherwise, if the switch resets, it will default to VTP server mode, and the extended-range VLAN IDs will not be saved.

To delete an extended-range VLAN, use the **no vlan *vlan-id*** global configuration command.

The procedure for assigning static-access ports to an extended-range VLAN is the same as for normal-range VLANs. See the “Assigning Static-Access Ports to a VLAN” section on page 251.

This example shows how to create a new extended-range VLAN with all default characteristics, enter config-vlan mode, and save the new VLAN in the switch startup configuration file:

```
Switch(config)# vtp mode transparent
Switch(config)# vlan 2000
Switch(config-vlan)# end
Switch# copy running-config startup config
```

Displaying VLANs

Use the **show vlan** privileged EXEC command to display a list of all VLANs on the switch, including extended-range VLANs. The display includes VLAN status, ports, and configuration information. To view normal-range VLANs in the VLAN database (1 to 1005) use the **show VLAN** configuration command (accessed by entering the **vlan database** privileged EXEC command). For a list of the VLAN IDs on the switch, use the **show running-config vlan** privileged EXEC command, optionally entering a VLAN ID range.

Table 49 lists the commands for monitoring VLANs.

Table 49. VLAN Monitoring Commands

Command	Command Mode	Purpose
show	VLAN configuration	Display status of VLANs in the VLAN database.
show current [<i>vlan-id</i>]	VLAN configuration	Display status of all or the specified VLAN in the VLAN database.
show interfaces [vlan <i>vlan-id</i>]	Privileged EXEC	Display characteristics for all interfaces or for the specified VLAN configured on the switch.
show running-config vlan	Privileged EXEC	Display all or a range of VLANs on the switch.
show vlan [id <i>vlan-id</i>]	Privileged EXEC	Display parameters for all VLANs or the specified VLAN on the switch.

For more details about the show command options and explanations of output fields, refer to the command reference for this release.

Configuring VLAN Trunks

These sections describe how VLAN trunks function on the switch:

- Trunking Overview, on page 255
- 802.1Q Configuration Considerations, on page 257
- Default Layer 2 Ethernet Interface VLAN Configuration, on page 257
- “Configuring an Ethernet Interface as a Trunk Port” section on page 258

Trunking Overview

A trunk is a point-to-point link between one or more Ethernet switch interfaces and another networking device such as a router or a switch. Gigabit Ethernet and Gigabit Ethernet trunks carry the traffic of multiple VLANs over a single link, and you can extend the VLANs across an entire network.

Figure 50 shows a network of switches that are connected by 802.1Q trunks.

Table 50. Layer 2 Interface Modes (continued)

Mode	Function
switchport mode dynamic auto	Makes the interface able to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to <i>trunk</i> or <i>desirable</i> mode.
switchport mode trunk	Puts the interface into permanent trunking mode and negotiates to convert the link into a trunk link. The interface becomes a trunk interface even if the neighboring interface is not a trunk interface.
switchport nonegotiate	Prevents the interface from generating DTP frames. You can use this command only when the interface switchport mode is access or trunk . You must manually configure the neighboring interface as a trunk interface to establish a trunk link.

802.1Q Configuration Considerations

802.1Q trunks impose these limitations on the trunking strategy for a network:

- In a network of Cisco switches connected through 802.1Q trunks, the switches maintain one instance of spanning tree for each VLAN allowed on the trunks. Non-Cisco devices might support one spanning-tree instance for all VLANs.

When you connect a Cisco switch to a non-Cisco device through an 802.1Q trunk, the Cisco switch combines the spanning-tree instance of the VLAN of the trunk with the spanning-tree instance of the non-Cisco 802.1Q switch. However, spanning-tree information for each VLAN is maintained by Cisco switches separated by a cloud of non-Cisco 802.1Q switches. The non-Cisco 802.1Q cloud separating the Cisco switches is treated as a single trunk link between the switches.

- Make sure the native VLAN for an 802.1Q trunk is the same on both ends of the trunk link. If the native VLAN on one end of the trunk is different from the native VLAN on the other end, spanning-tree loops might result.
- Disabling spanning tree on the native VLAN of an 802.1Q trunk without disabling spanning tree on every VLAN in the network can potentially cause spanning-tree loops. We recommend that you leave spanning tree enabled on the native VLAN of an 802.1Q trunk or disable spanning tree on every VLAN in the network. Make sure your network is loop-free before disabling spanning tree.

Default Layer 2 Ethernet Interface VLAN Configuration

Table 51 shows the default Layer 2 Ethernet interface VLAN configuration.

Table 51. Default Layer 2 Ethernet Interface VLAN Configuration

Feature	Default Setting
Allowed VLAN range	VLAN ID range is 2 to 4094 on the internal 1000 Mbps ports. VLAN ID range is 1 to 1006 on the internal 100 Mbps management module ports. VLAN ID range is 1 to 4094 on the external 10/100/1000 Mbps ports.
VLAN range eligible for pruning	VLANs 2 to 1001
Default VLAN (for external 10/100/1000 Mbps ports and internal 1000 Mbps ports)	VLAN 2

Table 51. Default Layer 2 Ethernet Interface VLAN Configuration (continued)

Feature	Default Setting
Default VLAN (for internal 100 Mbps management module ports)	VLAN 1
Native VLAN (for 802.1Q trunks)	VLAN 1 (for internal 100 Mbps management module ports) VLAN 2 (for external 10/100/1000 Mbps ports and internal 1000 Mbps ports)

Note: In trunk mode, the external ports will use VLAN 2 as the native VLAN. In access mode the external ports will use VLAN 1 as the default VLAN. If you are not using the management module ethernet interface, you can manage the switch using telnet by attaching to the switch directly through the external ports. This is possible because the default management VLAN for the switch is VLAN 1.

Configuring an Ethernet Interface as a Trunk Port

Because trunk ports send and receive VTP advertisements, to use VTP you must ensure that at least one trunk port is configured on the switch and that this trunk port is connected to the trunk port of a second switch. Otherwise, the switch cannot receive any VTP advertisements.

This section includes these procedures for configuring an Ethernet interface as a trunk port on the switch:

- Interaction with Other Features, on page 258
- “Configuring a Trunk Port” section on page 259
- Defining the Allowed VLANs on a Trunk, on page 260
- Changing the Pruning-Eligible List, on page 261
- Configuring the Native VLAN for Untagged Traffic, on page 262

Note: The default mode for the external 10/100/1000 Mbps interfaces is **switchport mode dynamic desirable**. If the neighboring interface supports trunking and is configured to allow trunking, the link is a Layer 2 trunk.

Interaction with Other Features

Trunking interacts with other features in these ways:

- A trunk port cannot be a secure port.
- Trunk ports can be grouped into EtherChannel port groups, but all trunks in the group must have the same configuration. When a group is first created, all ports follow the parameters set for the first port to be added to the group. If you change the configuration of one of these parameters, the switch propagates the setting that you entered to all ports in the group:
 - allowed-VLAN list
 - STP port priority for each VLAN
 - STP Port Fast setting
 - trunk status (If one port in a port group ceases to be a trunk, all ports cease to be trunks.)
- If you try to enable 802.1X on a trunk port, an error message appears, and 802.1X is not enabled. If you try to change the mode of an 802.1X-enabled port to trunk, the port mode is not changed.

- A port in dynamic mode can negotiate with its neighbor to become a trunk port. If you try to enable 802.1X on a dynamic port, an error message appears, and 802.1X is not enabled. If you try to change the mode of an 802.1X-enabled port to dynamic, the port mode is not changed.
- Protected ports are supported on 802.1Q trunks.

Configuring a Trunk Port

Beginning in privileged EXEC mode, follow these steps to configure a port as 802.1Q trunk port:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter the interface configuration mode and the port to be configured for trunking.
3.	switchport mode {dynamic {auto desirable} trunk}	Configure the interface as a Layer 2 trunk (required only if the interface is a Layer 2 access port or to specify the trunking mode). <ul style="list-style-type: none"> • dynamic auto—Set the interface to a trunk link if the neighboring interface is set to trunk or desirable mode. • dynamic desirable—Set the interface to a trunk link if the neighboring interface is set to trunk, desirable, or auto mode. • trunk—Set the interface in permanent trunking mode and negotiate to convert the link to a trunk link even if the neighboring interface is not a trunk interface.
4.	switchport access vlan <i>vlan-id</i>	(Optional) Specify the default VLAN, which is used if the interface stops trunking.
5.	switchport trunk native vlan <i>vlan-id</i>	Specify the native VLAN.
6.	end	Return to privileged EXEC mode.
7.	show interfaces <i>interface-id</i> switchport	Display the switchport configuration of the interface in the <i>Administrative Mode</i> and the <i>Administrative Trunking Encapsulation</i> fields of the display.
8.	show interfaces <i>interface-id</i> trunk	Display the trunk configuration of the interface.
9.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return an interface to its default configuration, use the **default interface** *interface-id* interface configuration command. To reset all trunking characteristics of a trunking interface to the defaults, use the **no switchport trunk** interface configuration command. To disable trunking, use the **switchport mode access** interface configuration command to configure the port as a static-access port.

This example shows how to configure the Gigabit Ethernet interface 0/17 as an 802.1Q trunk. The example assumes that the neighbor interface is configured to support 802.1Q trunking.

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# switchport mode dynamic desirable
Switch(config-if)# end
```

Defining the Allowed VLANs on a Trunk

Note: You cannot change the trunk mode on the internal interfaces connected to the 100 Mbps management module (ports 15 and 16). You also cannot remove the management VLAN from the allowed list.

By default, a trunk port sends traffic to and receives traffic from all VLANs. All VLAN IDs, 1 to 4094 are allowed on each trunk. However, you can remove VLANs from the allowed list, preventing traffic from those VLANs from passing over the trunk. To restrict the traffic a trunk carries, use the **switchport trunk allowed vlan remove *vlan-list*** interface configuration command to remove specific VLANs from the allowed list.

To reduce the risk of spanning-tree loops or storms, you can disable VLAN 1 on any individual VLAN trunk port by removing VLAN 1 from the allowed list. This is known as VLAN 1 minimization. VLAN 1 minimization disables VLAN 1 (the default VLAN on all Cisco switch trunk ports), on an individual VLAN trunk link. As a result no user traffic, including spanning-tree advertisements, are sent or received on VLAN 1.

When you remove VLAN 1 from a trunk port, the interface continues to send and receive management traffic, for example, Cisco Discovery Protocol (CDP), Port Aggregation Protocol (PAgP), Link Aggregation Control Protocol (LACP), Dynamic Trunking Protocol (DTP), and VLAN Trunking Protocol (VTP) in VLAN 1.

If a trunk port with VLAN 1 disabled is converted to a nontrunk port, it is added to the access VLAN. If the access VLAN is set to 1, then the port is added to VLAN 1, regardless of the **switchport trunk allowed** setting. The same is true for any VLAN that has been disabled on the port.

A trunk port can become a member of a VLAN if the VLAN is enabled, if VTP knows of the VLAN, and if the VLAN is in the allowed list for the port. When VTP detects a newly enabled VLAN and the VLAN is in the allowed list for a trunk port, the trunk port automatically becomes a member of the enabled VLAN. When VTP detects a new VLAN and the VLAN is not in the allowed list for a trunk port, the trunk port does not become a member of the new VLAN.

Beginning in privileged EXEC mode, follow these steps to modify the allowed list of an 802.1Q trunk:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode and the port to be configured.
3.	switchport mode trunk	Configure the interface as a VLAN trunk port.
4.	switchport trunk allowed vlan {add all except remove} <i>vlan-list</i>	(Optional) Configure the list of VLANs allowed on the trunk. For explanations about using the add , all , except , and remove keywords, refer to the command reference for this release. The <i>vlan-list</i> parameter is either a single VLAN number from 1 to 4094 or a range of VLANs described by two VLAN numbers, the lower one first, separated by a hyphen. Do not enter any spaces between comma-separated VLAN parameters or in hyphen-specified ranges. All VLANs are allowed by default.
5.	end	Return to privileged EXEC mode.

Step	Command	Purpose
6.	show interfaces <i>interface-id</i> switchport	Verify your entries in the <i>Trunking VLANs Enabled</i> field of the display.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default allowed VLAN list of all VLANs, use the **no switchport trunk allowed vlan** interface configuration command.

This example shows how to remove VLAN 2 from the allowed VLAN list:

```
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# switchport trunk allowed vlan remove 2
Switch(config-if)# end
Switch#
```

Changing the Pruning-Eligible List

The pruning-eligible list applies only to trunk ports. Each trunk port has its own eligibility list. VTP pruning must be enabled for this procedure to take effect. The “Enabling VTP Pruning” section on page 286 describes how to enable VTP pruning.

Beginning in privileged EXEC mode, follow these steps to remove VLANs from the pruning-eligible list on a trunk port:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and select the trunk port for which VLANs should be pruned.
3.	switchport trunk pruning vlan { add except none remove } <i>vlan-list</i> [, <i>vlan</i> [, <i>vlan</i> [,,,]]	Configure the list of VLANs allowed to be pruned from the trunk. (See the “VTP Pruning” section on page 278). For explanations about using the add , except , none , and remove keywords, refer to the command reference for this release. Separate nonconsecutive VLAN IDs with a comma and no spaces; use a hyphen to designate a range of IDs. Valid IDs are from 2 to 1001. Note: Extended-range VLANs (VLAN IDs 1006 to 4094) cannot be pruned. VLANs that are pruning-ineligible receive flooded traffic. The default list of VLANs allowed to be pruned contains VLANs 2 to 1001.
4.	end	Return to privileged EXEC mode.
5.	show interfaces <i>interface-id</i> switchport	Verify your entries in the <i>Pruning VLANs Enabled</i> field of the display.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default pruning-eligible list of all VLANs, use the **no switchport trunk pruning vlan** interface configuration command.

Configuring the Native VLAN for Untagged Traffic

A trunk port configured with 802.1Q tagging can receive both tagged and untagged traffic. By default, the switch forwards untagged traffic in the native VLAN configured for the port. The native VLAN is VLAN 1 by default.

Note: The native VLAN can be assigned any VLAN ID; it is not dependent on the management VLAN.

Note: The native VLAN ID on the internal 100 Mbps management module interfaces (ports 15 and 16) changes when the management VLAN changes. The native VLAN cannot be explicitly changed, it will only change when the management VLAN of the switch changes. Changing the native VLAN on management module interfaces is not allowed. This ensures that the switch and the management module always have an open communication path for ethernet traffic used to manage the switch.

For information about 802.1Q configuration issues, see the “802.1Q Configuration Considerations” section on page 257.

Beginning in privileged EXEC mode, follow these steps to configure the native VLAN on an 802.1Q trunk:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and define the interface that is configured as the 802.1Q trunk.
3.	switchport trunk native vlan <i>vlan-id</i>	Configure the VLAN that is sending and receiving untagged traffic on the trunk port. For <i>vlan-id</i> , the range is 1 to 4094.
4.	end	Return to privileged EXEC mode.
5.	show interfaces <i>interface-id</i> switchport	Verify your entries in the <i>Trunking Native Mode VLAN</i> field.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default native VLAN, VLAN 1, use the **no switchport trunk native vlan** interface configuration command.

If a packet has a VLAN ID that is the same as the outgoing port native VLAN ID, the packet is sent untagged; otherwise, the switch sends the packet with a tag.

Load Sharing Using STP

Load sharing divides the bandwidth supplied by parallel trunks connecting switches. To avoid loops, STP normally blocks all but one parallel link between switches. Using load sharing, you divide the traffic between the links according to which VLAN the traffic belongs.

You configure load sharing on trunk ports by using STP port priorities or STP path costs. For load sharing using STP port priorities, both load-sharing links must be connected to the same switch. For load sharing using STP path costs, each load-sharing link can be connected to the same switch or to two different switches. For more information about STP, see Chapter 10 “Configuring STP.”

Load Sharing Using STP Port Priorities


When two ports on the same switch form a loop, the STP port priority setting determines which port is enabled and which port is in a blocking state. You can set the priorities on a parallel trunk port so that the port carries all the traffic for a given VLAN. The trunk port with the higher priority (lower values) for a VLAN is forwarding traffic for that VLAN. The trunk port with the lower priority (higher values) for the same VLAN remains in a blocking state for that VLAN. One trunk port sends or receives all traffic for the VLAN.

Figure 51 shows two trunks connecting supported switches. In this example, the switches are configured as follows:

- VLANs 8 through 10 are assigned a port priority of 10 on Trunk 1.
- VLANs 3 through 6 retain the default port priority of 128 on Trunk 1.
- VLANs 3 through 6 are assigned a port priority of 10 on Trunk 2.
- VLANs 8 through 10 retain the default port priority of 128 on Trunk 2.

In this way, Trunk 1 carries traffic for VLANs 8 through 10, and Trunk 2 carries traffic for VLANs 3 through 6. If the active trunk fails, the trunk with the lower priority takes over and carries the traffic for all of the VLANs. No duplication of traffic occurs over any trunk port.

Figure 51. Load Sharing by Using STP Port Priorities



Beginning in privileged EXEC mode, follow these steps to configure the network shown in Figure 51.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode on Switch 1.
2.	vtp domain <i>domain-name</i>	Configure a VTP administrative domain. The domain name can be from 1 to 32 characters.
3.	vtp mode server	Configure Switch 1 as the VTP server.
4.	end	Return to privileged EXEC mode.
5.	show vtp status	Verify the VTP configuration on both Switch 1 and Switch 2. In the display, check the <i>VTP Operating Mode</i> and the <i>VTP Domain Name</i> fields.
6.	show vlan	Verify that the VLANs exist in the database on Switch 1.
7.	configure terminal	Enter global configuration mode.
8.	interface gigabitethernet 0/17	Enter interface configuration mode, and define Gigabit Ethernet port 0/17 as the interface to be configured as a trunk.

Step	Command	Purpose
9.	switchport mode trunk	Configure the port as a trunk port.
10.	end	Return to privilege EXEC mode.
11.	show interfaces gigabitethernet0/18 switchport	Verify the VLAN configuration.
12.		Repeat Steps 7 through 11 on Switch 1 for Gigabit Ethernet port 0/18.
13.		Repeat Steps 7 through 11 on Switch 2 to configure the trunk ports on Gigabit Ethernet ports 0/17 and 0/18.
14.	show vlan	When the trunk links come up, VTP passes the VTP and VLAN information to Switch 2. Verify that Switch 2 has learned the VLAN configuration.
15.	configure terminal	Enter global configuration mode on Switch 1.
16.	interface gigabitethernet0/17	Enter interface configuration mode, and define the interface to set the STP port priority.
17.	spanning-tree vlan 8 port-priority 10	Assign the port priority of 10 for VLAN 8.
18.	spanning-tree vlan 9 port-priority 10	Assign the port priority of 10 for VLAN 9.
19.	spanning-tree vlan 10 port-priority 10	Assign the port priority of 10 for VLAN 10.
20.	exit	Return to global configuration mode.
21.	interface gigabitethernet0/18	Enter interface configuration mode, and define the interface to set the STP port priority.
22.	spanning-tree vlan 3 port-priority 10	Assign the port priority of 10 for VLAN 3.
23.	spanning-tree vlan 4 port-priority 10	Assign the port priority of 10 for VLAN 4.
24.	spanning-tree vlan 5 port-priority 10	Assign the port priority of 10 for VLAN 5.
25.	spanning-tree vlan 6 port-priority 10	Assign the port priority of 10 for VLAN 6.
26.	end	Return to privileged EXEC mode.
27.	show running-config	Verify your entries.
28.	copy running-config startup-config	(Optional) Save your entries in the configuration file.


Load Sharing Using STP Path Cost

You can configure parallel trunks to share VLAN traffic by setting different path costs on a trunk and associating the path costs with different sets of VLANs. The VLANs keep the traffic separate. Because no loops exist, STP does not disable the ports, and redundancy is maintained in the event of a lost link.

In Figure 52, Trunk ports 1 and 2 are 1000BASE-T ports. The path costs for the VLANs are assigned as follows:

- VLANs 2 through 4 are assigned a path cost of 30 on Trunk port 1.
- VLANs 8 through 10 retain the default 1000BASE-T path cost on Trunk port 1 of 4.
- VLANs 8 through 10 are assigned a path cost of 30 on Trunk port 2.
- VLANs 2 through 4 retain the default 1000BASE-T path cost on Trunk port 2 of 4.

Figure 52. Load-Sharing Trunks with Traffic Distributed by Path Cost



16591

Beginning in privileged EXEC mode, follow these steps to configure the network shown in Figure 52:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode on Switch 1.
2.	interface gigabitethernet 0/17	Enter interface configuration mode, and define Gigabit Ethernet port 0/17 as the interface to be configured as a trunk.
3.	switchport mode trunk	Configure the port as a trunk port.
4.	exit	Return to global configuration mode.
5.		Repeat Steps 2 through 4 on Switch 1 interface Gigabit Ethernet 0/18.
6.	end	Return to privileged EXEC mode.
7.	show running-config	Verify your entries. In the display, make sure that interfaces Gigabit Ethernet 0/17 and Gigabit Ethernet 0/18 are configured as trunk ports.
8.	show vlan	When the trunk links come up, Switch 1 receives the VTP information from the other switches. Verify that Switch 1 has learned the VLAN configuration.
9.	configure terminal	Enter global configuration mode.
10.	interface gigabitethernet 0/17	Enter interface configuration mode, and define Gigabit Ethernet port 0/17 as the interface to set the STP cost.
11.	spanning-tree vlan 2 cost 30	Set the spanning-tree path cost to 30 for VLAN 2.
12.	spanning-tree vlan 3 cost 30	Set the spanning-tree path cost to 30 for VLAN 3.
13.	spanning-tree vlan 4 cost 30	Set the spanning-tree path cost to 30 for VLAN 4.
14.	end	Return to global configuration mode.
15.		Repeat Steps 9 through 11 on Switch 1 interface Gigabit Ethernet 0/18, and set the spanning-tree path cost to 30 for VLANs 8, 9, and 10.
16.	exit	Return to privileged EXEC mode.

Step	Command	Purpose
17.	show running-config	Verify your entries. In the display, verify that the path costs are set correctly for interfaces Gigabit Ethernet 0/17 and 0/18.
18.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Configuring VMPS

The switch cannot be a VMPS server but can act as a client to the VMPS and communicate with it through the VLAN Query Protocol (VQP). VMPS dynamically assigns dynamic access port VLAN membership.

This section includes this information about configuring VMPS:

- “Understanding VMPS” section on page 266
- “Default VMPS Configuration” section on page 269
- “VMPS Configuration Guidelines” section on page 269
- “Configuring the VMPS Client” section on page 270
- “Monitoring the VMPS” section on page 272
- “Troubleshooting Dynamic Port VLAN Membership” section on page 273
- “VMPS Configuration Example” section on page 273

Understanding VMPS

When the VMPS receives a VQP request from a client switch, it searches its database for a MAC-address-to-VLAN mapping. The server response is based on this mapping and whether or not the server is in secure mode. Secure mode determines whether the server shuts down the port when a VLAN is not allowed on it or just denies the port access to the VLAN.

In response to a request, the VMPS takes one of these actions:

- If the assigned VLAN is restricted to a group of ports, the VMPS verifies the requesting port against this group and responds as follows:
 - If the VLAN is allowed on the port, the VMPS sends the VLAN name to the client in response.
 - If the VLAN is not allowed on the port and the VMPS is not in secure mode, the VMPS sends an *access-denied* response.
 - If the VLAN is not allowed on the port and the VMPS is in secure mode, the VMPS sends a *port-shutdown* response.
- If the VLAN in the database does not match the current VLAN on the port and active hosts exist on the port, the VMPS sends an *access-denied* or a *port-shutdown* response, depending on the secure mode of the VMPS.

If the switch receives an *access-denied* response from the VMPS, it continues to block traffic from the MAC address to or from the port. The switch continues to monitor the packets directed to the port and sends a query to the VMPS when it identifies a new address. If the switch receives a *port-shutdown* response from the VMPS, it disables the port. The port must be manually re-enabled by using the CLI, CMS, or SNMP.

You can also use an explicit entry in the configuration table to deny access to specific MAC addresses for security reasons. If you enter the **none** keyword for the VLAN

name, the VMPS sends an *access-denied* or *port-shutdown* response, depending on the VMPS secure mode setting.

Dynamic Port VLAN Membership

A dynamic (nontrunking) port on the switch can belong to only one VLAN, with a VLAN ID from 1 to 1005. When the link comes up, the switch does not forward traffic to or from this port until the VMPS provides the VLAN assignment. The VMPS receives the source MAC address from the first packet of a new host connected to the dynamic port and attempts to match the MAC address to a VLAN in the VMPS database.

If there is a match, the VMPS sends the VLAN number for that port. If the client switch was not previously configured, it uses the domain name from the first VTP packet it receives on its trunk port from the VMPS. If the client switch was previously configured, it includes its domain name in the query packet to the VMPS to obtain its VLAN number. The VMPS verifies that the domain name in the packet matches its own domain name before accepting the request and responds to the client with the assigned VLAN number for the client. If there is no match, the VMPS either denies the request or shuts down the port (depending on the VMPS secure mode setting).

Multiple hosts (MAC addresses) can be active on a dynamic port if they are all in the same VLAN; however, the VMPS shuts down a dynamic port if more than 20 hosts are active on the port.

If the link goes down on a dynamic port, the port returns to an isolated state and does not belong to a VLAN. Any hosts that come online through the port are checked again through the VQP with the VMPS before the port is assigned to a VLAN.

VMPS Database Configuration File

The VMPS contains a database configuration file that you create. This ASCII text file is stored on a switch-accessible TFTP server that functions as a server for VMPS. The file contains VMPS information, such as the domain name, the fallback VLAN name, and the MAC-address-to-VLAN mapping. The switch cannot act as the VMPS, but you can use a Catalyst 5000 or Catalyst 6000 series switch as the VMPS.

You can configure a fallback VLAN name. If you connect a device with a MAC address that is not in the database, the VMPS sends the fallback VLAN name to the client. If you do not configure a fallback VLAN and the MAC address does not exist in the database, the VMPS sends an *access-denied* response. If the VMPS is in secure mode, it sends a *port-shutdown* response.

Whenever port names are used in the VMPS database configuration file, the server must use the switch convention for naming ports. For example, Gi0/17 is fixed Gigabit Ethernet port number 17. If the switch is a cluster member, the command switch adds the name of the switch before the type. For example, *es3%Gi0/17* refers to fixed Gigabit Ethernet port 17 on member switch 3. When port names are required, these naming conventions must be followed in the VMPS database configuration file when it is configured to support a cluster.

This example shows a example of a VMPS database configuration file as it appears on a Catalyst 6000 series switch. The file has these characteristics:

- The security mode is open.
- The default is used for the fallback VLAN.
- MAC address-to-VLAN name mappings—The MAC address of each host and the VLAN to which each host belongs is defined.
- Port groups are defined.
- VLAN groups are defined.

- VLAN port policies are defined for the ports associated with restricted VLANs.

```

!VMPS File Format, version 1.1
! Always begin the configuration file with
! the word "VMPS"
!
!vmps domain <domain-name>
! The VMPS domain must be defined.
!vmps mode {open | secure}
! The default mode is open.
!vmps fallback <vlan-name>
!vmps no-domain-req { allow | deny }
!
! The default value is allow.
vmps domain DSBU
vmps mode open
vmps fallback default
vmps no-domain-req deny
!
!
!MAC Addresses
!
vmps-mac-addr
!
! address <addr> vlan-name <vlan_name>
!
address 0012.2233.4455 vlan-name hardware
address 0000.6509.a080 vlan-name hardware
address aabb.ccdd.eeff vlan-name Green
address 1223.5678.9abc vlan-name ExecStaff
address fedc.ba98.7654 vlan-name --NONE--
address fedc.ba23.1245 vlan-name Purple
!
!Port Groups
!
!vmps-port-group <group-name>
! device <device-id> { port <port-name> | all-ports }
!
vmps-port-group WiringCloset1
  device 198.92.30.32 port 0/2
  device 172.20.26.141 port 0/8
vmps-port-group "Executive Row"
  device 198.4.254.222 port 0/2
  device 198.4.254.222 port 0/3
  device 198.4.254.223 all-ports
!
!
!VLAN groups
!
!vmps-vlan-group <group-name>
! vlan-name <vlan-name>
!
vmps-vlan-group Engineering
  vlan-name hardware
  vlan-name software
!
!
!VLAN port Policies
!

```

```

!vmmps-port-policies {vlan-name <vlan_name> | vlan-group <group-name> }
! { port-group <group-name> | device <device-id> port <port-name> }
!
vmmps-port-policies vlan-group Engineering
port-group WiringCloset1
vmmps-port-policies vlan-name Green
device 198.92.30.32 port 0/8
vmmps-port-policies vlan-name Purple
device 198.4.254.22 port 0/2
port-group "Executive Row"

```

Default VMPS Configuration

Table 52 shows the default VMPS and dynamic port configuration on client switches.

Table 52. Default VMPS Client and Dynamic Port Configuration

Feature	Default Setting
VMPS domain server	None
VMPS reconfirm interval	60 minutes
VMPS server retry count	3
Dynamic ports	None configured

VMPS Configuration Guidelines

These guidelines and restrictions apply to dynamic access port VLAN membership:

- You should configure the VMPS before you configure ports as dynamic.
- The communication between a cluster of switches and VMPS is managed by the command switch and includes port-naming conventions that are different from standard port names. For the cluster-based port-naming conventions, see the “VMPS Database Configuration File” section on page 267.
- When you configure a port as a dynamic access port, the spanning-tree Port Fast feature is automatically enabled for that port. The Port Fast mode accelerates the process of bringing the port into the forwarding state.
- 802.1X ports cannot be configured as dynamic access ports. If you try to enable 802.1X on a dynamic-access (VQP) port, an error message appears, and 802.1X is not enabled. If you try to change an 802.1X-enabled port to dynamic VLAN assignment, an error message appears, and the VLAN configuration is not changed.
- Trunk ports cannot be dynamic access ports, but you can enter the **switchport access vlan dynamic** interface configuration command for a trunk port. In this case, the switch retains the setting and applies it if the port is later configured as an access port.
You must turn off trunking on the port before the dynamic access setting takes effect.
- Dynamic access ports cannot be network ports or monitor ports.
- Secure ports cannot be dynamic access ports. You must disable port security on a port before it becomes dynamic.
- Dynamic access ports cannot be members of an EtherChannel group.
- Port channels cannot be configured as dynamic access ports.
- The VTP management domain of the VMPS client and the VMPS server must be the same.

- VQP does not support extended-range VLANs (VLAN IDs higher than 1006). Extended-range VLANs cannot be configured by VMPS.
- The VLAN configured on the VMPS server should not be a voice VLAN.

Configuring the VMPS Client

You configure dynamic VLANs by using the VMPS (server). The switch can be a VMPS client; it cannot be a VMPS server.

Entering the IP Address of the VMPS

You must first enter the IP address of the server to configure the switch as a client.

Note: If the VMPS is being defined for a cluster of switches, enter the address on the command switch.

Beginning in privileged EXEC mode, follow these steps to enter the IP address of the VMPS:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	vmps server <i>ipaddress</i> primary	Enter the IP address of the switch acting as the primary VMPS server.
3.	vmps server <i>ipaddress</i>	Enter the IP address of the switch acting as a secondary VMPS server. You can enter up to three secondary server addresses.
4.	end	Return to privileged EXEC mode.
5.	show vmps	Verify your entries in the <i>VMPS Domain Server</i> field of the display.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Note: The switch port that is connected to the VMPS server cannot be a dynamic access port. It can be either a static access port or a trunk port. See the “Configuring an Ethernet Interface as a Trunk Port” section on page 258.

Configuring Dynamic Access Ports on VMPS Clients

If you are configuring a port on a cluster member switch as a dynamic port, first use the **rcommand** privileged EXEC command to log into the member switch.

Caution: Dynamic port VLAN membership is for end stations or hubs connected to end stations. Connecting dynamic access ports to other switches can cause a loss of connectivity.

Beginning in privileged EXEC mode, follow these steps to configure a dynamic access port on a VMPS client switch:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode and the switch port that is connected to the end station.
3.	switchport mode access	Set the port to access mode.

Step	Command	Purpose
4.	switchport access vlan dynamic	Configure the port as eligible for dynamic VLAN membership. The dynamic access port must be connected to an end station.
5.	end	Return to privileged EXEC mode.
6.	show interfaces <i>interface-id</i> switchport	Verify your entries in the <i>Operational Mode</i> field of the display.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return an interface to its default configuration, use the **default interface** *interface-id* interface configuration command. To return an interface to its default switchport mode (dynamic desirable), use the **no switchport mode** interface configuration command. To reset the access mode to the default VLAN for the switch, use the **no switchport access** interface configuration command.

Note: When you configure a dynamic access port by using the **switchport access vlan dynamic** interface configuration command, the port might allow unauthorized users to access network resources if the interface changes from access mode to trunk mode through the DTP negotiation. The workaround is to configure the port as a static access port.

Reconfirming VLAN Memberships

Beginning in privileged EXEC mode, follow these steps to confirm the dynamic port VLAN membership assignments that the switch has received from the VMPS:

Step	Command	Purpose
1.	vmps reconfirm	Reconfirm dynamic port VLAN membership.
2.	show vmps	Verify the dynamic VLAN reconfirmation status.

Changing the Reconfirmation Interval

VMPS clients periodically reconfirm the VLAN membership information received from the VMPS. You can set the number of minutes after which reconfirmation occurs.

If you are configuring a member switch in a cluster, this parameter must be equal to or greater than the reconfirmation setting on the command switch. You must also first use the **rcommand** privileged EXEC command to log into the member switch.

Beginning in privileged EXEC mode, follow these steps to change the reconfirmation interval:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	vmps reconfirm <i>minutes</i>	Enter the number of minutes between reconfirmations of the dynamic VLAN membership. Enter a number from 1 to 120. The default is 60 minutes.
3.	end	Return to privileged EXEC mode.
4.	show vmps	Verify the dynamic VLAN reconfirmation status in the <i>Reconfirm Interval</i> field of the display.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no vmps reconfirm** global configuration command.

Changing the Retry Count

Beginning in privileged EXEC mode, follow these steps to change the number of times that the switch attempts to contact the VMPS before querying the next server:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	vmps retry count	Change the retry count. The retry range is from 1 to 10; the default is 3.
3.	end	Return to privileged EXEC mode.
4.	show vmps	Verify your entry in the <i>Server Retry Count</i> field of the display.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no vmps retry** global configuration command.

Monitoring the VMPS

You can display information about the VMPS by using the **show vmps** privileged EXEC command. The switch displays this information about the VMPS:

VMPS VQP Version	The version of VQP used to communicate with the VMPS. The switch queries the VMPS that is using VQP version 1.
Reconfirm Interval	The number of minutes the switch waits before reconfirming the VLAN-to-MAC-address assignments.
Server Retry Count	The number of times VQP resends a query to the VMPS. If no response is received after this many tries, the switch starts to query the secondary VMPS.
VMPS domain server	The IP address of the configured VLAN membership policy servers. The switch sends queries to the one marked <i>current</i> . The one marked <i>primary</i> is the primary server.
VMPS Action	The result of the most recent reconfirmation attempt. A reconfirmation attempt can occur automatically when the reconfirmation interval expired, or you can force it by entering the vmps reconfirm privileged EXEC command or its CMS or SNMP equivalent.

This is an example of output for the **show vmps** privileged EXEC command:

```
Switch# show vmps

VQP Client Status:
-----
VMPS VQP Version: 1
Reconfirm Interval: 60 min
Server Retry Count: 3
VMPS domain server: 172.20.128.86 (primary, current)
                   172.20.128.87

Reconfirmation status
```

VMPS Action: No Dynamic Port

Troubleshooting Dynamic Port VLAN Membership

The VMPS shuts down a dynamic port under these conditions:

- The VMPS is in secure mode, and it does not allow the host to connect to the port. The VMPS shuts down the port to prevent the host from connecting to the network.
- More than 20 active hosts reside on a dynamic port.

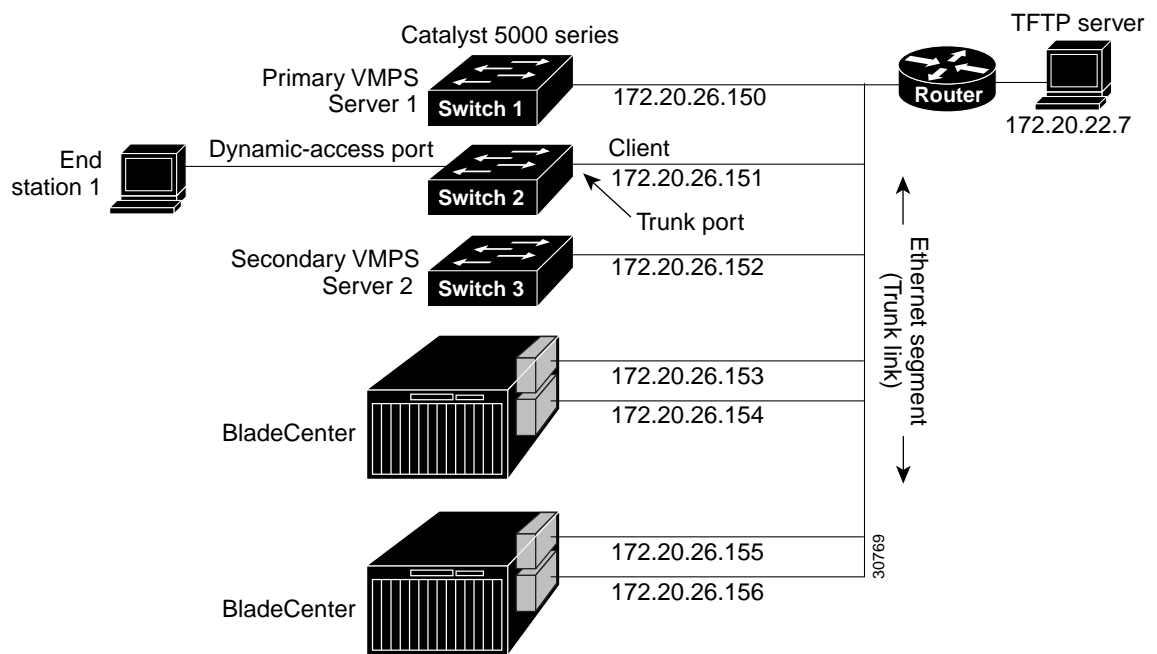
To re-enable a disabled dynamic port, enter the **no shutdown** interface configuration command.

VMPS Configuration Example

Figure 53 shows a network with a VMPS server switch and VMPS client switches with dynamic ports. In this example, these assumptions apply:

- The VMPS server and the VMPS client are separate switches.
- The Catalyst 5000 series Switch 1 is the primary VMPS server.
- The Catalyst 5000 series Switch 3 and Switch 10 are secondary VMPS servers.
- The end stations are connected to these clients:
 - Catalyst 2950 Switch 2
 - Catalyst 3500 XL Switch 9
- The database configuration file is stored on the TFTP server with the IP address 172.20.22.7.

Figure 53. Dynamic Port VLAN Membership Configuration



Chapter 14. Configuring VTP

This chapter describes how to use the VLAN Trunking Protocol (VTP) and the VLAN database for managing VLANs on your switch.

Note: For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

The chapter includes these sections:

- Understanding VTP, on page 275
- Configuring VTP, on page 280
- Monitoring VTP, on page 288

Understanding VTP

VTP is a Layer 2 messaging protocol that maintains VLAN configuration consistency by managing the addition, deletion, and renaming of VLANs on a network-wide basis. VTP minimizes misconfigurations and configuration inconsistencies that can cause several problems, such as duplicate VLAN names, incorrect VLAN-type specifications, and security violations.

Before you create VLANs, you must decide whether to use VTP in your network. Using VTP, you can make configuration changes centrally on one or more switches and have those changes automatically communicated to all the other switches in the network. Without VTP, you cannot send information about VLANs to other switches. VTP configuration information is saved in the VTP VLAN database.

VTP only learns about normal-range VLANs (VLAN IDs 1 to 1005). Extended-range VLANs (VLAN IDs greater than 1005) are not supported by VTP or stored in the VTP VLAN database.

This section contains information about these VTP parameters:

- The VTP Domain, on page 275
- VTP Modes, on page 276
- VTP Advertisements, on page 277
- VTP Version 2, on page 278
- VTP Pruning, on page 278

The VTP Domain

A VTP domain (also called a VLAN management domain) consists of one switch or several interconnected switches under the same administrative responsibility sharing the same VTP domain name. A switch can be in only one VTP domain. You make global VLAN configuration changes for the domain by using the command-line interface (CLI), Cluster Management Suite (CMS) software, or Simple Network Management Protocol (SNMP).

By default, the switch is in VTP no-management-domain state until it receives an advertisement for a domain over a trunk link (a link that carries the traffic of multiple VLANs) or until you configure a domain name. Until the management domain name is specified or learned, you cannot create or modify VLANs on a VTP server, and VLAN information is not propagated over the network.

If the switch receives a VTP advertisement over a trunk link, it inherits the management domain name and the VTP configuration revision number. The switch

then ignores advertisements with a different domain name or an earlier configuration revision number.

Caution: Before adding a VTP client switch to a VTP domain, always verify that its VTP configuration revision number is *lower* than the configuration revision number of the other switches in the VTP domain. Switches in a VTP domain always use the VLAN configuration of the switch with the highest VTP configuration revision number. If you add a switch that has a revision number higher than the revision number in the VTP domain, it can erase all VLAN information from the VTP server and VTP domain. See the “Adding a VTP Client Switch to a VTP Domain” section on page 287 for the procedure for verifying and resetting the VTP configuration revision number.

When you make a change to the VLAN configuration on a VTP server, the change is propagated to all switches in the VTP domain. VTP advertisements are sent over all IEEE 802.1Q trunk connections. VTP maps VLANs dynamically across multiple LAN types with unique names and internal index associates. Mapping eliminates excessive device administration required from network administrators.

If you configure a switch for VTP transparent mode, you can create and modify VLANs, but the changes are not sent to other switches in the domain, and they affect only the individual switch. However, configuration changes made when the switch is in this mode are saved in the switch running configuration and can be saved to the switch startup configuration file.

For domain name and password configuration guidelines, see the “VTP Configuration Guidelines” section on page 281.

VTP Modes

You can configure a supported switch to be in one of the VTP modes listed in Table 53.

Table 53. VTP Modes

VTP Mode	Description
VTP server	<p>In VTP server mode, you can create, modify, and delete VLANs and specify other configuration parameters (such as the VTP version) for the entire VTP domain. VTP servers advertise their VLAN configurations to other switches in the same VTP domain and synchronize their VLAN configurations with other switches based on advertisements received over trunk links.</p> <p>In VTP server mode, VLAN configurations are saved in nonvolatile RAM (NVRAM).</p>

Table 53. VTP Modes (continued)

VTP Mode	Description
VTP client	<p>A VTP client behaves like a VTP server, but you cannot create, change, or delete VLANs on a VTP client.</p> <p>In VTP client mode, VLAN configurations are not saved in NVRAM.</p>
VTP transparent	<p>VTP transparent switches do not participate in VTP. A VTP transparent switch does not advertise its VLAN configuration and does not synchronize its VLAN configuration based on received advertisements. However, in VTP version 2, transparent switches do forward VTP advertisements that they receive from other switches from their trunk interfaces. You can create, modify, and delete VLANs on a switch in VTP transparent mode. The switch must be in VTP transparent mode when you create extended-range VLANs. See the “Configuring Extended-Range VLANs” section on page 252.</p> <p>When the switch is in VTP transparent mode, the VTP and VLAN configurations are saved in NVRAM, but they are not advertised to other switches. In this mode, VTP mode and domain name are saved in the switch running configuration and you can save this information in the switch startup configuration file by entering the copy running-config startup-config privileged EXEC command.</p>

When the network is configured with more than the maximum 250 VLANs, the switch automatically changes from VTP server or client mode to VTP transparent mode. The switch then operates with the VLAN configuration that preceded the one that sent it into transparent mode.

VTP Advertisements

Each switch in the VTP domain sends periodic global configuration advertisements from each trunk port to a reserved multicast address. Neighboring switches receive these advertisements and update their VTP and VLAN configurations as necessary.

Note: Because trunk ports send and receive VTP advertisements, you must ensure that at least one trunk port is configured on the switch and that this trunk port is connected to the trunk port of a second switch. Otherwise, the switch cannot receive any VTP advertisements. For more information on trunk ports, see the “Configuring VLAN Trunks” section on page 255.

VTP advertisements distribute this global domain information:

- VTP domain name
- VTP configuration revision number
- Update identity and update timestamp
- MD5 digest VLAN configuration, including maximum transmission unit (MTU) size for each VLAN.
- Frame format

VTP advertisements distribute this VLAN information for each configured VLAN:

- VLAN IDs
- VLAN name
- VLAN type
- VLAN state
- Additional VLAN configuration information specific to the VLAN type

VTP Version 2

If you use VTP in your network, you must decide whether to use version 1 or version 2. By default, VTP operates in version 1.

VTP version 2 supports these features not supported in version 1:

- Token Ring support—VTP version 2 supports Token Ring Bridge Relay Function (TrBRF) and Token Ring Concentrator Relay Function (TrCRF) VLANs. For more information about Token Ring VLANs, see the “Configuring Normal-Range VLANs” section on page 245.
- Unrecognized Type-Length-Value (TLV) support—A VTP server or client propagates configuration changes to its other trunks, even for TLVs it is not able to parse. The unrecognized TLV is saved in NVRAM when the switch is operating in VTP server mode.
- Version-Dependent Transparent Mode—In VTP version 1, a VTP transparent switch inspects VTP messages for the domain name and version and forwards a message only if the version and domain name match. Because VTP version 2 supports only one domain, it forwards VTP messages in transparent mode without inspecting the version and domain name.
- Consistency Checks—In VTP version 2, VLAN consistency checks (such as VLAN names and values) are performed only when you enter new information through the CLI, the Cluster Management Software (CMS), or SNMP. Consistency checks are not performed when new information is obtained from a VTP message or when information is read from NVRAM. If the MD5 digest on a received VTP message is correct, its information is accepted.

VTP Pruning

VTP pruning increases network available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to reach the destination devices. Without VTP pruning, a switch floods broadcast, multicast, and unknown unicast traffic across all trunk links within a VTP domain even though receiving switches might discard them. VTP pruning is disabled by default.

VTP pruning blocks unneeded flooded traffic to VLANs on trunk ports that are included in the pruning-eligible list. Only VLANs included in the pruning-eligible list can be pruned. By default, VLANs 2 through 1001 are pruning eligible switch trunk ports. If the VLANs are configured as pruning-ineligible, the flooding continues. VTP pruning is supported with VTP version 1 and version 2.

Figure 54 shows a switched network without VTP pruning enabled. Port 1 on Switch 1 and Port 2 on Switch 4 are assigned to the Red VLAN. If a broadcast is sent from the host connected to Switch 1, Switch 1 floods the broadcast and every switch in the network receives it, even though Switches 3, 5, and 6 have no ports in the Red VLAN.

Figure 54. Flooding Traffic without VTP Pruning





Figure 55 shows a switched network with VTP pruning enabled. The broadcast traffic from Switch 1 is not forwarded to Switches 3, 5, and 6 because traffic for the Red VLAN has been pruned on the links shown (Port 5 on Switch 2 and Port 4 on Switch 4).

Figure 55. Optimized Flooded Traffic with VTP Pruning



Enabling VTP pruning on a VTP server enables pruning for the entire management domain. Making VLANs pruning-eligible or pruning-ineligible affects pruning eligibility for those VLANs on that device only (not on all switches in the VTP domain). See the “Enabling VTP Pruning” section on page 286. VTP pruning takes effect several seconds after you enable it. VTP pruning does not prune traffic from VLANs that are pruning-ineligible. VLAN 1 and VLANs 1002 to 1005 are always pruning-ineligible; traffic from these VLANs cannot be pruned. Extended-range VLANs (VLAN IDs higher than 1005) are also pruning-ineligible.

VTP pruning is not designed to function in VTP transparent mode. If one or more switches in the network are in VTP transparent mode, you should do one of these:

- Turn off VTP pruning in the entire network.
- Turn off VTP pruning by making all VLANs on the trunk of the switch upstream to the VTP transparent switch pruning ineligible.

To configure VTP pruning on an interface, use the **switchport trunk pruning vlan** interface configuration command (see the “Changing the Pruning-Eligible List” section on page 261). VTP pruning operates when an interface is trunking. You can set VLAN pruning-eligibility, whether or not VTP pruning is enabled for the VTP domain, whether or not any given VLAN exists, and whether or not the interface is currently trunking.

Configuring VTP

This section includes guidelines and procedures for configuring VTP. These sections are included:

- Default VTP Configuration, on page 280
- VTP Configuration Options, on page 280
- VTP Configuration Guidelines, on page 281
- Configuring a VTP Server, on page 282
- Configuring a VTP Client, on page 284
- Disabling VTP (VTP Transparent Mode), on page 285
- Enabling VTP Version 2, on page 286
- Enabling VTP Pruning, on page 286
- Adding a VTP Client Switch to a VTP Domain, on page 287

Default VTP Configuration

Table 54 shows the default VTP configuration.

Table 54. Default VTP Configuration

Feature	Default Setting
VTP domain name	Null.
VTP mode	Transparent
VTP version 2 enable state	Version 2 is disabled.
VTP password	None.
VTP pruning	Disabled.

VTP Configuration Options

You can configure VTP by using these configuration modes.

- VTP Configuration in Global Configuration Mode, on page 280
- VTP Configuration in VLAN Configuration Mode, on page 281

You access VLAN configuration mode by entering the **vlan database** privileged EXEC command.

For detailed information about **vtp** commands, refer to the command reference for this release.

VTP Configuration in Global Configuration Mode

You can use the **vtp** global configuration command to set the VTP password, the version, the VTP file name, the interface providing updated VTP information, the domain name, and the mode, and to disable or enable pruning. For more information

about available keywords, refer to the command descriptions in the command reference for this release. The VTP information is saved in the VTP VLAN database. When VTP mode is transparent, the VTP domain name and mode are also saved in the switch running configuration file, and you can save it in the switch startup configuration file by entering the **copy running-config startup-config** privileged EXEC command. You must use this command if you want to save VTP mode as transparent, even if the switch resets.

When you save VTP information in the switch startup configuration file and reboot the switch, the switch configuration is determined as follows:

- If the VTP mode is transparent in the startup configuration and the VLAN database and the VTP domain name from the VLAN database matches that in the startup configuration file, the VLAN database is ignored (cleared), and the VTP and VLAN configurations in the startup configuration file are used. The VLAN database revision number remains unchanged in the VLAN database.
- If the VTP mode or domain name in the startup configuration do not match the VLAN database, the domain name and VTP mode and configuration for the first 1005 VLANs use the VLAN database information.

VTP Configuration in VLAN Configuration Mode

You can configure all VTP parameters in VLAN configuration mode, which you access by entering the **vlan database** privileged EXEC command. For more information about available keywords, refer to the **vtp** VLAN configuration command description in the command reference for this release. When you enter the **exit** command in VLAN configuration mode, it applies all the commands that you entered and updates the VLAN database. VTP messages are sent to other switches in the VTP domain, and the privileged EXEC mode prompt appears.

If VTP mode is transparent, the domain name and the mode (transparent) are saved in the switch running configuration, and you can save this information in the switch startup configuration file by entering the **copy running-config startup-config** privileged EXEC command.

VTP Configuration Guidelines

These sections describe guidelines you should follow when implementing VTP in your network.

Domain Names

When configuring VTP for the first time, you must always assign a domain name. You must configure all switches in the VTP domain with the same domain name. Switches in VTP transparent mode do not exchange VTP messages with other switches, and you do not need to configure a VTP domain name for them.

Note: If NVRAM and DRAM storage is sufficient, all switches in a VTP domain should be in VTP server mode.

Caution: Do not configure a VTP domain if all switches are operating in VTP client mode. If you configure the domain, it is impossible to make changes to the VLAN configuration of that domain. Make sure that you configure at least one switch in the VTP domain for VTP server mode.

Passwords

You can configure a password for the VTP domain, but it is not required. If you do configure a domain password, all domain switches must share the same password

and you must configure the password on each switch in the management domain. Switches without a password or with the wrong password reject VTP advertisements.

If you configure a VTP password for a domain, a switch that is booted without a VTP configuration does not accept VTP advertisements until you configure it with the correct password. After the configuration, the switch accepts the next VTP advertisement that uses the same password and domain name in the advertisement.

If you are adding a new switch to an existing network with VTP capability, the new switch learns the domain name only after the applicable password has been configured on it.

Caution: When you configure a VTP domain password, the management domain does not function properly if you do not assign a management domain password to each switch in the domain.

VTP Version

Follow these guidelines when deciding which VTP version to implement:

- All switches in a VTP domain must run the same VTP version.
- A VTP version 2-capable switch can operate in the same VTP domain as a switch running VTP version 1 if version 2 is disabled on the version 2-capable switch (version 2 is disabled by default).
- Do not enable VTP version 2 on a switch unless all of the switches in the same VTP domain are version-2-capable. When you enable version 2 on a switch, all of the version-2-capable switches in the domain enable version 2. If there is a version 1-only switch, it does not exchange VTP information with switches with version 2 enabled.
- If there are TrBRF and TrCRF Token Ring networks in your environment, you must enable VTP version 2 for Token Ring VLAN switching to function properly. To run Token Ring and Token Ring-Net, disable VTP version 2.

Configuration Requirements

When you configure VTP, you must configure a trunk port so that the switch can send and receive VTP advertisements. For more information, see the “Configuring VLAN Trunks” section on page 255.

If you are configuring VTP on a cluster member switch to a VLAN, use the **rcommand** privileged EXEC command to log into the member switch. For more information about the command, refer to the command reference for this release.

If you are configuring extended-range VLANs on the switch, the switch must be in VTP transparent mode.

Configuring a VTP Server

When a switch is in VTP server mode, you can change the VLAN configuration and have it propagated throughout the network.

Note: If extended-range VLANs are configured on the switch, you cannot change VTP mode to server. You receive an error message, and the configuration is not allowed.

Beginning in privileged EXEC mode, follow these steps to configure the switch as a VTP server:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	vtp mode server	Configure the switch for VTP server mode.
3.	vtp domain <i>domain-name</i>	Configure the VTP administrative-domain name. The name can be from 1 to 32 characters. All switches operating in VTP server or client mode under the same administrative responsibility must be configured with the same domain name.
4.	vtp password <i>password</i>	(Optional) Set the password for the VTP domain. The password can be from 8 to 64 characters. If you configure a VTP password, the VTP domain does not function properly if you do not assign the same password to each switch in the domain.
5.	end	Return to privileged EXEC mode.
6.	show vtp status	Verify your entries in the <i>VTP Operating Mode</i> and the <i>VTP Domain Name</i> fields of the display.

When you configure a domain name, it cannot be removed; you can only reassign a switch to a different domain.

To return the switch to a no-password state, use the **no vtp password** global configuration command.

This example shows how to use global configuration mode to configure the switch as a VTP server with the domain name *eng_group* and the password *mypassword*:

```
Switch# config terminal
Switch(config)# vtp mode server
Switch(config)# vtp domain eng_group
Switch(config)# vtp password mypassword
Switch(config)# end
```

You can also use VLAN configuration mode to configure VTP parameters.

Beginning in privileged EXEC mode, follow these steps to use VLAN configuration mode to configure the switch as a VTP server:

Step	Command	Purpose
1.	vlan database	Enter VLAN configuration mode.
2.	vtp server	Configure the switch for VTP server mode.
3.	vtp domain <i>domain-name</i>	Configure a VTP administrative-domain name. The name can be from 1 to 32 characters. All switches operating in VTP server or client mode under the same administrative responsibility must be configured with the same domain name.
4.	vtp password <i>password</i>	(Optional) Set a password for the VTP domain. The password can be from 8 to 64 characters. If you configure a VTP password, the VTP domain does not function properly if you do not assign the same password to each switch in the domain.

Step	Command	Purpose
5.	exit	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
6.	show vtp status	Verify your entries in the <i>VTP Operating Mode</i> and the <i>VTP Domain Name</i> fields of the display.

When you configure a domain name, it cannot be removed; you can only reassign a switch to a different domain.

To return the switch to a no-password state, use the **no vtp password** VLAN configuration command.

This example shows how to use VLAN configuration mode to configure the switch as a VTP server with the domain name *eng_group* and the password *mypassword*:

```
Switch# vlan database
Switch(vlan)# vtp server
Switch(vlan)# vtp domain eng_group
Switch(vlan)# vtp password mypassword
Switch(vlan)# exit
APPLY completed.
Exiting....
```

Configuring a VTP Client

When a switch is in VTP client mode, you cannot change its VLAN configuration. The client switch receives VTP updates from a VTP server in the VTP domain and then modifies its configuration accordingly.

Note: If extended-range VLANs are configured on the switch, you cannot change VTP mode to client. You receive an error message, and the configuration is not allowed.

Caution: If all switches are operating in VTP client mode, do not configure a VTP domain name. If you do, it is impossible to make changes to the VLAN configuration of that domain. Therefore, make sure you configure at least one switch as a VTP server.

Beginning in privileged EXEC mode, follow these steps to configure the switch as a VTP client:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	vtp mode client	Configure the switch for VTP client mode.
3.	vtp domain <i>domain-name</i>	(Optional) Enter the VTP administrative-domain name. The name can be from 1 to 32 characters. This should be the same domain name as the VTP server. All switches operating in VTP server or client mode under the same administrative responsibility must be configured with the same domain name.
4.	vtp password <i>password</i>	(Optional) Enter the password for the VTP domain.
5.	end	Return to privileged EXEC mode.
6.	show vtp status	Verify your entries in the <i>VTP Operating Mode</i> and the <i>VTP Domain Name</i> fields of the display.

Use the **no vtp mode** global configuration command to return the switch to VTP server mode. To return the switch to a no-password state, use the **no vtp password** global configuration command. When you configure a domain name, it cannot be removed; you can only reassign a switch to a different domain.

Note: You can also configure a VTP client by using the **vlan database** privileged EXEC command to enter VLAN configuration mode and entering the **vtp client** command, similar to the second procedure under “Configuring a VTP Server” section on page 282. Use the **no vtp client** VLAN configuration command to return the switch to VTP server mode or the **no vtp password** VLAN configuration command to return the switch to a no-password state. When you configure a domain name, it cannot be removed; you can only reassign a switch to a different domain.

Disabling VTP (VTP Transparent Mode)

When you configure the switch for VTP transparent mode, you disable VTP on the switch. The switch does not send VTP updates and does not act on VTP updates received from other switches. However, a VTP transparent switch running VTP version 2 does forward received VTP advertisements on all of its trunk links.

Note: Before you create extended-range VLANs (VLAN IDs 1006 to 4094), you must set VTP mode to transparent by using the **vtp mode transparent** global configuration command. Save this configuration to the startup configuration so that the switch boots up in VTP transparent mode. Otherwise, you lose the extended-range VLAN configuration if the switch resets and boots up in VTP server mode.

Beginning in privileged EXEC mode, follow these steps to configure VTP transparent mode and save the VTP configuration in the switch startup configuration file:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	vtp mode transparent	Configure the switch for VTP transparent mode (disable VTP).
3.	end	Return to privileged EXEC mode.
4.	show vtp status	Verify your entries in the <i>VTP Operating Mode</i> and the <i>VTP Domain Name</i> fields of the display.
5.	copy running-config startup-config	(Optional) Save the configuration in the startup configuration file. Note: Only VTP mode and domain name are saved in the switch running configuration and can be copied to the startup configuration file.

To return the switch to VTP server mode, use the **no vtp mode** global configuration command.

Note: If extended-range VLANs are configured on the switch, you cannot change VTP mode to server. You receive an error message, and the configuration is not allowed.

Note: You can also configure VTP transparent mode by using the **vlan database** privileged EXEC command to enter VLAN configuration mode and by entering the **vtp transparent** command, similar to the second procedure under the “Configuring a VTP Server” section on page 282. Use the **no vtp transparent** VLAN configuration command to return the switch to VTP server mode. If extended-range VLANs are configured on the switch, you cannot change VTP mode to server. You receive an error message, and the configuration is not allowed.

Enabling VTP Version 2

VTP version 2 is disabled by default on VTP version 2-capable switches. When you enable VTP version 2 on a switch, every VTP version 2-capable switch in the VTP domain enables version 2. You can only configure the version on switches in VTP server or transparent mode.

Caution: VTP version 1 and VTP version 2 are not interoperable on switches in the same VTP domain. Every switch in the VTP domain must use the same VTP version. Do not enable VTP version 2 unless every switch in the VTP domain supports version 2.

Note: In TrCRF and TrBRF Token ring environments, you must enable VTP version 2 for Token Ring VLAN switching to function properly. For Token Ring and Token Ring-Net media, VTP version 2 must be disabled.

For more information on VTP version configuration guidelines, see the “VTP Version” section on page 282.

Beginning in privileged EXEC mode, follow these steps to enable VTP version 2:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	vtp version 2	Enable VTP version 2 on the switch. VTP version 2 is disabled by default on VTP version 2-capable switches.
3.	end	Return to privileged EXEC mode.
4.	show vtp status	Verify that VTP version 2 is enabled in the <i>VTP V2 Mode</i> field of the display.

To disable VTP version 2, use the **no vtp version** global configuration command.

Note: You can also enable VTP version 2 by using the **vlan database** privileged EXEC command to enter VLAN configuration mode and entering the **vtp v2-mode** VLAN configuration command. To disable VTP version 2, use the **no vtp v2-mode** VLAN configuration command.

Enabling VTP Pruning

Pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the destination devices. You can only enable VTP pruning on a switch in VTP server mode.

Beginning in privileged EXEC mode, follow these steps to enable VTP pruning in the VTP domain:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	vtp pruning	Enable pruning in the VTP administrative domain. By default, pruning is disabled. You need to enable pruning on only one switch in VTP server mode.
3.	end	Return to privileged EXEC mode.
4.	show vtp status	Verify your entries in the <i>VTP Pruning Mode</i> field of the display.

To disable VTP pruning, use the **no vtp pruning** global configuration command.

Note: You can also enable VTP pruning by using the **vlan database** privileged EXEC command to enter VLAN configuration mode and entering the **vtp pruning** VLAN configuration command. To disable VTP pruning, use the **no vtp pruning** VLAN configuration command.

Pruning is supported with VTP version 1 and version 2. If you enable pruning on the VTP server, it is enabled for the entire VTP domain.

Only VLANs included in the pruning-eligible list can be pruned. By default, VLANs 2 through 1001 are pruning eligible on trunk ports. Reserved VLANs and extended-range VLANs cannot be pruned. To change the pruning-eligible VLANs, see the "Changing the Pruning-Eligible List" section on page 261.

Adding a VTP Client Switch to a VTP Domain

Before adding a VTP client to a VTP domain, always verify that its VTP configuration revision number is **lower** than the configuration revision number of the other switches in the VTP domain. Switches in a VTP domain always use the VLAN configuration of the switch with the highest VTP configuration revision number. If you add a switch that has a revision number higher than the revision number in the VTP domain, it can erase all VLAN information from the VTP server and VTP domain.

Beginning in privileged EXEC mode, follow these steps to verify and reset the VTP configuration revision number on a switch *before* adding it to a VTP domain:

Step	Command	Purpose
1.	show vtp status	Check the VTP configuration revision number. If the number is 0, add the switch to the VTP domain. If the number is greater than 0, follow these steps: a. Write down the domain name. b. Write down the configuration revision number. c. Continue with the next steps to reset the configuration revision number on the switch.
2.	configure terminal	Enter global configuration mode.
3.	vtp domain <i>domain-name</i>	Change the domain name from the original one displayed in Step 1 to a new name.
4.	end	The VLAN information on the switch is updated and the configuration revision number is reset to 0. You return to privileged EXEC mode.
5.	show vtp status	Verify that the configuration revision number has been reset to 0.
6.	configure terminal	Enter global configuration mode.
7.	vtp domain <i>domain-name</i>	Enter the original domain name on the switch.
8.	end	The VLAN information on the switch is updated, and you return to privileged EXEC mode.
9.	show vtp status	(Optional) Verify that the domain name is the same as in Step 1 and that the configuration revision number is 0.

You can also change the VTP domain name by entering the **vlan database** privileged EXEC command to enter VLAN configuration mode and by entering the **vtp domain** *domain-name* command. In this mode, you must enter the **exit** command to update VLAN information and return to privileged EXEC mode.

After resetting the configuration revision number, add the switch to the VTP domain.

Note: You can use the **vtp mode transparent** global configuration command or the **vtp transparent** VLAN configuration command to disable VTP on the switch, and then change its VLAN information without affecting the other switches in the VTP domain.

Monitoring VTP

You monitor VTP by displaying VTP configuration information: the domain name, the current VTP revision, and the number of VLANs. You can also display statistics about the advertisements sent and received by the switch.

Table 55 shows the privileged EXEC commands for monitoring VTP activity.

Table 55. VTP Monitoring Commands

Command	Purpose
show vtp status	Display the VTP switch configuration information.
show vtp counters	Display counters about VTP messages that have been sent and received.

This is an example of output from the **show vtp status** privileged EXEC command:

```
Switch# show vtp status
VTP Version                : 2
Configuration Revision     : 25
Maximum VLANs supported locally : 250
Number of existing VLANs   : 69
VTP Operating Mode         : Server
VTP Domain Name            : test
VTP Pruning Mode           : Disabled
VTP V2 Mode                : Disabled
VTP Traps Generation       : Disabled
MD5 digest                 : 0x59 0xBA 0x92 0xA4 0x74 0xD5 0x42 0x29
Configuration last modified by 0.0.0.0 at 3-1-93 00:18:42
Local updater ID is 10.1.1.59 on interface V11 (lowest numbered VLAN interface found)
```

This is an example of output from the **show vtp counters** privileged EXEC command:

```
Switch# show vtp counters
VTP statistics:
Summary advertisements received : 20
Subset advertisements received  : 0
Request advertisements received  : 0
Summary advertisements transmitted : 11
Subset advertisements transmitted : 0
Request advertisements transmitted : 0
Number of config revision errors  : 0
Number of config digest errors    : 0
Number of V1 summary errors       : 0

VTP pruning statistics:

Trunk          Join Transmitted Join Received  Summary advts received from
-----          -----          -----          -----
non-pruning-capable device
```

Chapter 15. Configuring IGMP Snooping and MVR

This chapter describes how to configure Internet Group Management Protocol (IGMP) snooping on your switch, including an application of local IGMP snooping, Multicast VLAN Registration (MVR). It also includes procedures for controlling multicast group membership by using IGMP filtering.

Note: For complete syntax and usage information for the commands used in this chapter, refer to the switch command reference for this release and the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference*.

This chapter consists of these sections:

- Understanding IGMP Snooping, on page 289
- Configuring IGMP Snooping, on page 293
- Displaying IGMP Snooping Information, on page 299
- Understanding Multicast VLAN Registration, on page 301
- Configuring MVR, on page 303
- Displaying MVR Information, on page 307
- Configuring IGMP Filtering, on page 308
- Displaying IGMP Filtering Configuration, on page 312

Note: For MAC addresses that map to IP multicast groups, you can either manage them through features such as IGMP snooping and MVR, or you can use static MAC addresses. However, you cannot use both methods simultaneously. Therefore, before using IGMP snooping or MVR, you should remove all statically configured MAC addresses that map to IP multicast groups.

Understanding IGMP Snooping

Layer 2 switches can use IGMP snooping to constrain the flooding of multicast traffic by dynamically configuring Layer 2 interfaces so that multicast traffic is forwarded to only those interfaces associated with IP multicast devices. As the name implies, IGMP snooping requires the LAN switch to snoop on the IGMP transmissions between the host and the router and to keep track of multicast groups and member ports. When the switch receives an IGMP report from a host for a particular multicast group, the switch adds the host port number to the forwarding table entry; when it receives an IGMP Leave Group message from a host, it removes the host port from the table entry. It also periodically deletes entries if it does not receive IGMP membership reports from the multicast clients.

Note: For more information on IP multicast and IGMP, refer to RFC 1112 and RFC 2236.

The multicast router sends out periodic IGMP general queries to all VLANs. When IGMP snooping is enabled, the switch responds to the router queries with only one join request per MAC multicast group, and the switch creates one entry per VLAN in the Layer 2 forwarding table for each MAC group from which it receives an IGMP join request. All hosts interested in this multicast traffic send join requests and are added to the forwarding table entry.

Layer 2 multicast groups learned through IGMP snooping are dynamic. However, you can statically configure MAC multicast groups by using the **ip igmp snooping vlan static** global configuration command. If you specify group membership for a multicast group address statically, your setting supersedes any automatic manipulation by

IGMP snooping. Multicast group membership lists can consist of both user-defined and IGMP snooping-learned settings.


If a port spanning-tree, a port group, or a VLAN ID change occurs, the IGMP snooping-learned multicast groups from this port on the VLAN are deleted.

The switches support a maximum of 255 IP multicast groups and support both IGMP version 1 and IGMP version 2.

Joining a Multicast Group

When a host connected to the switch wants to join an IP multicast group, it sends an unsolicited IGMP join message, specifying the IP multicast group to join. Alternatively, when the switch receives a general query from the router, it forwards the query to all ports in the VLAN. Hosts wanting to join the multicast group respond by sending a join message to the switch. The switch CPU creates a multicast forwarding-table entry for the group if it is not already present. The CPU also adds the interface where the join message was received to the forwarding-table entry. The host associated with that interface receives multicast traffic for that multicast group. See Figure 56.

Figure 56. Initial IGMP Join Message



Router A sends a general query to the switch, which forwards the query to ports 2 through 5, all members of the same VLAN. Host 1 wants to join multicast group 224.1.2.3 and multicasts an IGMP membership report (IGMP join message) to the group with the equivalent MAC destination address of 0x0100.5E01.0203. When the CPU receives the IGMP report multicast by Host 1, the CPU uses the information in the IGMP report to set up a forwarding-table entry, as shown in Table 56, that includes the port numbers of Host 1, the router, and the switch internal CPU.

Table 56. IGMP Snooping Forwarding Table

Destination Address	Type of Packet	Ports
0100.5exx.xxxx	IGMP	0
0100.5e01.0203	!IGMP	1, 2

Note that the switch hardware can distinguish IGMP information packets from other packets for the multicast group.

- The first entry in the table tells the switching engine to send IGMP packets to only the switch CPU. This prevents the CPU from becoming overloaded with multicast frames.
- The second entry tells the switching engine to send frames addressed to the 0x0100.5E01.0203 multicast MAC address that are not IGMP packets (!IGMP) to the router and to the host that has joined the group.

If another host (for example, Host 4) sends an unsolicited IGMP join message for the same group (Figure 57), the CPU receives that message and adds the port number of Host 4 to the forwarding table as shown in Table 57. Note that because the forwarding table directs IGMP messages to only the CPU, the message is not flooded to other ports on the switch. Any known multicast traffic is forwarded to the group and not to the CPU. Any unknown multicast traffic is flooded to the VLAN and sent to the CPU until it becomes known.

Figure 57. Second Host Joining a Multicast Group

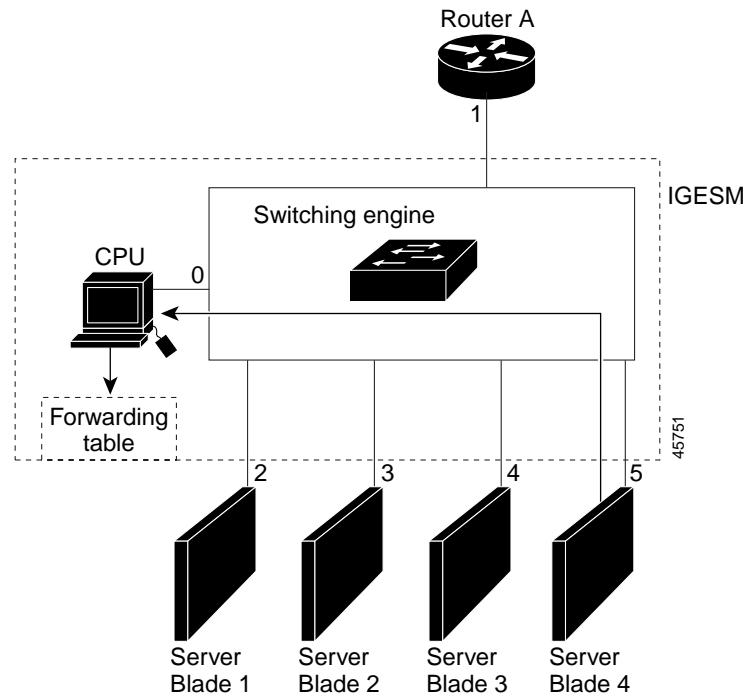


Table 57. Updated IGMP Snooping Forwarding Table

Destination Address	Type of Packet	Ports
0100.5exx.xxxx	IGMP	0
0100.5e01.0203	!IGMP	1, 2, 5

Leaving a Multicast Group

The router sends periodic multicast general queries and the switch forwards these queries through all ports in the VLAN. Interested hosts respond to the queries. If at least one host in the VLAN wishes to receive multicast traffic, the router continues forwarding the multicast traffic to the VLAN. The switch forwards multicast group traffic to only those hosts listed in the forwarding table for that Layer 2 multicast group.

When hosts want to leave a multicast group, they can either silently leave, or they can send a leave message. When the switch receives a leave message from a host, it sends out a MAC-based general query to determine if any other devices connected to that interface are interested in traffic for the specific multicast group. The switch then updates the forwarding table for that MAC group so that only those hosts interested in receiving multicast traffic for the group are listed in the forwarding table. If the router receives no reports from a VLAN, it removes the group for the VLAN from its IGMP cache.

Immediate-Leave Processing

The switch uses IGMP snooping Immediate-Leave processing to remove from the forwarding table an interface that sends a leave message without the switch sending MAC-based general queries to the interface. The VLAN interface is pruned from the multicast tree for the multicast group specified in the original leave message. Immediate-Leave processing ensures optimal bandwidth management for all hosts on a switched network, even when multiple multicast groups are simultaneously in use.

Note: You should only use the Immediate-Leave processing feature on VLANs where a single host is connected to each port. If Immediate Leave is enabled in VLANs where more than one host is connected to a port, some hosts might be inadvertently dropped. Immediate Leave is supported with only IGMP version 2 hosts.

Source-Only Networks

In a source-only network, switch ports are connected to multicast source ports and multicast router ports. The switch ports are not connected to hosts that send IGMP join or leave messages.

The switch learns about IP multicast groups from the IP multicast data stream by using the source-only learning method. The switch forwards traffic only to the multicast router ports.

The default learning method is IP multicast-source-only learning. You can disable IP multicast-source-only learning by using the **no ip igmp snooping source-only-learning** global configuration command.

By default, the switch ages out forwarding-table entries that were learned by the source-only learning method and that are not in use. If the aging time is too long or is disabled, the forwarding table is filled with unused entries that the switch learned by using source-only learning or by using the IGMP join messages. When the switch receives traffic for new IP multicast groups, it floods the packet to all ports in the same VLAN. This unnecessary flooding can impact switch performance.

If aging is disabled and you want to delete multicast addresses that the switch learned by using source-only learning, re-enable aging of the forwarding-table entries. The switch can now age out the multicast addresses that were learned by the source-only learning method and that are not in use.

Configuring IGMP Snooping

IGMP snooping allows switches to examine IGMP packets and make forwarding decisions based on their content.

These sections describe how to configure IGMP snooping:

- Default IGMP Snooping Configuration, on page 293
- Enabling or Disabling IGMP Snooping, on page 293
- Setting the Snooping Method, on page 294
- Configuring a Multicast Router Port, on page 295
- Configuring a Host Statically to Join a Group, on page 296
- Enabling IGMP Immediate-Leave Processing, on page 297
- Disabling IP Multicast-Source-Only Learning, on page 297
- Configuring the Aging Time, on page 298

Default IGMP Snooping Configuration

Table 58 shows the default IGMP snooping configuration.

Table 58. Default IGMP Snooping Configuration

Feature	Default Setting
IGMP snooping	Enabled globally and per VLAN
Multicast routers	None configured
Multicast router learning (snooping) method	PIM-DVMRP
IGMP snooping Immediate Leave	Disabled
Static groups	None configured
IP multicast-source-only learning	Enabled
Aging forward-table entries (when source-only learning is enabled)	Enabled. The default is 600 seconds (10 minutes).

Enabling or Disabling IGMP Snooping

By default, IGMP snooping is globally enabled on the switch. When globally enabled or disabled, it is also enabled or disabled in all existing VLAN interfaces. IGMP snooping is by default enabled on all VLANs, but can be enabled and disabled on a per-VLAN basis.

Global IGMP snooping overrides the VLAN IGMP snooping. If global snooping is disabled, you cannot enable VLAN snooping. If global snooping is enabled, you can enable or disable VLAN snooping.

Beginning in privileged EXEC mode, follow these steps to globally enable IGMP snooping on the switch:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	ip igmp snooping	Globally enable IGMP snooping in all existing VLAN interfaces.

Step	Command	Purpose
3.	end	Return to privileged EXEC mode.
4.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To globally disable IGMP snooping on all VLAN interfaces, use the **no ip igmp snooping** global configuration command.

Beginning in privileged EXEC mode, follow these steps to enable IGMP snooping on a VLAN interface:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	ip igmp snooping vlan <i>vlan-id</i>	Enable IGMP snooping on the VLAN interface.
3.	end	Return to privileged EXEC mode.
4.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable IGMP snooping on a VLAN interface, use the **no ip igmp snooping vlan *vlan-id*** global configuration command for the specified VLAN number.

Setting the Snooping Method

Multicast-capable router ports are added to the forwarding table for every Layer 2 multicast entry. The switch learns of such ports through one of these methods:

- Snooping on IGMP queries, Protocol Independent Multicast (PIM) packets, and Distance Vector Multicast Routing Protocol (DVMRP) packets
- Listening to Cisco Group Management Protocol (CGMP) packets from other routers
- Statically connecting to a multicast router port with the **ip igmp snooping mrouter** global configuration command

You can configure the switch either to snoop on IGMP queries and PIM/DVMRP packets or to listen to CGMP self-join or proxy-join packets. By default, the switch snoops on PIM/DVMRP packets on all VLANs. To learn of multicast router ports through only CGMP packets, use the **ip igmp snooping vlan *vlan-id* mrouter learn cgmp** global configuration command. When this command is entered, the router listens to only CGMP self-join and CGMP proxy-join packets and no other CGMP packets. To learn of multicast router ports through only PIM-DVMRP packets, use the **ip igmp snooping vlan *vlan-id* mrouter learn pim-dvmrp** global configuration command.

Beginning in privileged EXEC mode, follow these steps to alter the method in which a VLAN interface dynamically accesses a multicast router:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	ip igmp snooping vlan <i>vlan-id</i> mrouter learn {cgmp pim-dvmrp}	Enable IGMP snooping on a VLAN. The VLAN ID range is 1 to 4094. Specify the multicast router learning method: <ul style="list-style-type: none"> • cgmp—Listen for CGMP packets. This method is useful for reducing control traffic. • pim-dvmrp—Snoop on IGMP queries and PIM-DVMRP packets. This is the default.
3.	end	Return to privileged EXEC mode.
4.	show ip igmp snooping	Verify the configuration.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to configure IGMP snooping to use CGMP packets as the learning method:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 1 mrouter learn cgmp
Switch(config)# end
Switch# show ip igmp snooping vlan 1
vlan 1
-----
IGMP snooping is globally enabled
IGMP snooping is enabled on this Vlan
IGMP snooping immediate-leave is disabled on this Vlan
IGMP snooping mrouter learn mode is cgmp on this Vlan
IGMP snooping is running in IGMP_ONLY mode on this Vlan
```

To return to the default learning method, use the **no ip igmp snooping vlan *vlan-id* mrouter learn cgmp** global configuration command.

Configuring a Multicast Router Port

To add a multicast router port (add a static connection to a multicast router), use the **ip igmp snooping vlan mrouter** global configuration command on the switch.

Beginning in privileged EXEC mode, follow these steps to enable a static connection to a multicast router:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	ip igmp snooping vlan <i>vlan-id</i> mrouter interface <i>interface-id</i>	Specify the multicast router VLAN ID and specify the interface to the multicast router. For the VLAN ID, the range is 1 to 4094.
3.	end	Return to privileged EXEC mode.
4.	show ip igmp snooping mrouter [vlan <i>vlan-id</i>]	Verify that IGMP snooping is enabled on the VLAN interface.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a multicast router port from the VLAN, use the **no ip igmp snooping vlan *vlan-id* mrouter interface *interface-id*** global configuration command.

This example shows how to enable a static connection to a multicast router and verify the configuration:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 200 mrouter interface
gigabitethernet0/17
Switch(config)# end
Switch# show ip igmp snooping mrouter vlan 200
vlan          ports
-----+-----
200          Gi0/17(static)
```

Configuring a Host Statically to Join a Group

Hosts or Layer 2 ports normally join multicast groups dynamically, but you can also statically configure a host on an interface.

Beginning in privileged EXEC mode, follow these steps to add a Layer 2 port as a member of a multicast group:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode
2.	ip igmp snooping vlan <i>vlan-id</i> static <i>mac-address</i> interface <i>interface-id</i>	Statically configure a Layer 2 port as a member of a multicast group: <ul style="list-style-type: none"> <i>vlan-id</i> is the multicast group VLAN ID. <i>mac-address</i> is the group MAC address. <i>interface-id</i> is the member port.
3.	end	Return to privileged EXEC mode.
4.	show ip igmp snooping mrouter vlan <i>vlan-id</i> or show mac address-table multicast vlan <i>vlan-id</i>	Verify that the member port is a member of the VLAN multicast group. Verify the member port and the MAC address
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the Layer 2 port from the multicast group, use the **no ip igmp snooping vlan *vlan-id* static *mac-address* interface *interface-id*** global configuration command.

This example shows how to statically configure a host on an interface and verify the configuration:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 1 static 0100.5e00.0203 interface
gigabitethernet0/17
Switch(config)# end
Switch# show mac address-table multicast vlan 1
Vlan    Mac Address          Type    Ports
----    -
1       0100.5e00.0203      USER   Gi0/17
```

Enabling IGMP Immediate-Leave Processing

When you enable IGMP Immediate-Leave processing, the switch immediately removes a port when it detects an IGMP version 2 leave message on that port. You should use the Immediate-Leave feature only when there is a single receiver present on every port in the VLAN.

Immediate Leave is supported with only IGMP version 2 hosts.

Beginning in privileged EXEC mode, follow these steps to enable IGMP Immediate-Leave processing:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode
2.	ip igmp snooping vlan <i>vlan-id</i> immediate-leave	Enable IGMP Immediate-Leave processing on the VLAN interface.
3.	end	Return to privileged EXEC mode.
4.	show ip igmp snooping vlan <i>vlan-id</i>	Verify that Immediate Leave is enabled on the VLAN.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable IGMP Immediate-Leave on a VLAN, use the **no ip igmp snooping vlan *vlan-id* immediate-leave** global configuration command.

This example shows how to enable IGMP immediate-leave processing on VLAN 130 and verify the configuration:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 130 immediate-leave
Switch(config)# end
Switch# show ip igmp snooping vlan 130
vlan 130
-----
IGMP snooping is globally enabled
IGMP snooping is enabled on this Vlan
IGMP snooping immediate-leave is enabled on this Vlan
IGMP snooping mrouter learn mode is pim-dvmrp on this Vlan
IGMP snooping is running in IGMP_ONLY mode on this Vlan
```

Disabling IP Multicast-Source-Only Learning

The IP multicast-source-only learning method is enabled by default. The switch learns the IP multicast group from the IP multicast data stream and only forwards traffic to the multicast router ports.

If IP multicast-source-only learning is disabled by using the **no ip igmp snooping source-only-learning** global configuration command, the switch floods unknown multicast traffic to the VLAN and sends the traffic to the CPU until the traffic becomes known. When the switch receives an IGMP report from a host for a particular multicast group, the switch forwards traffic from this multicast group only to the multicast router ports.

Note: Do not disable IP multicast-source-only learning. IP multicast-source-only learning should be disabled only if your network is not composed of IP multicast-source-only networks and if disabling this learning method improves the network performance.

Beginning in privileged EXEC mode, follow these steps to disable IP multicast-source-only learning:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode
2.	no ip igmp snooping source-only-learning	Disable IP multicast-source-only learning.
3.	end	Return to privileged EXEC mode.
4.	show running-config include source-only-learning	Verify that IP multicast-source-only learning is disabled.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To enable IP multicast-source-only learning, use the **ip igmp snooping source-only-learning** global configuration command.

This example shows how to disable IP multicast-source-only learning and verify the configuration:

```
Switch# configure terminal
Switch(config)# no ip igmp snooping source-only-learning
Switch(config)# end
Switch# show running-config | include source-only-learning
Current configuration : 1972 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname Switch
!
enable password my_password
!
ip subnet-zero
no ip igmp snooping source-only-learning
!
!
spanning-tree extend system-id
no spanning-tree vlan 1
!
!
interface gigabitethernet0/17
no ip address
!
<output truncated>
```

Configuring the Aging Time

You can set the aging time for forwarding-table entries that the switch learns by using the IP multicast-source-only learning method.

Beginning in privileged EXEC mode, follow these steps to configure the aging time:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode
2.	ip igmp snooping source-only-learning age-timer <i>time</i>	Set the aging time. The range is from 0 to 2880 seconds. The default is 600 seconds (10 minutes).
3.	end	Return to privileged EXEC mode.
4.	show running-config include source-only-learning	Verify that the aging time.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the aging of the forwarding table entries, enter the **ip igmp snooping source-only-learning age-timer 0** global configuration command.

If you disable source-only learning by using the **no ip igmp snooping source-only learning** global configuration command and the aging time is enabled, it has no effect on the switch.

Displaying IGMP Snooping Information

You can display IGMP snooping information for dynamically learned and statically configured router ports and VLAN interfaces. You can also display MAC address multicast entries for a VLAN configured for IGMP snooping.

To display IGMP snooping information, use one or more of the privileged EXEC commands in Table 59.

Table 59. Commands for Displaying IGMP Snooping Information

Command	Purpose
show ip igmp snooping [vlan <i>vlan-id</i>]	Display the snooping configuration information for all VLANs on the switch or for a specified VLAN. (Optional) Enter vlan <i>vlan-id</i> to display information for a single VLAN.
show ip igmp snooping mrouter [vlan <i>vlan-id</i>]	Display information on dynamically learned and manually configured multicast router interfaces. Note: When you enable IGMP snooping, the switch automatically learns the interface to which a multicast router is connected. These are dynamically learned interfaces. (Optional) Enter vlan <i>vlan-id</i> to display information for a single VLAN.
show mac address-table multicast [vlan <i>vlan-id</i>] [user igmp-snooping] [count]	Display the Layer 2 MAC address table entries for a VLAN. The keywords are all optional and limit the display as shown: <ul style="list-style-type: none"> • vlan <i>vlan-id</i>—Displays only the specified multicast group VLAN. • user—Displays only the user-configured multicast entries. • igmp-snooping—Displays only entries learned through IGMP snooping. • count—Displays only the total number of entries for the selected criteria, not the actual entries.

This is an example of output from the **show ip igmp snooping** privileged EXEC command for all VLAN interfaces on the switch:

```
Switch# show ip igmp snooping
vlan 1
-----
IGMP snooping is globally enabled
IGMP snooping is enabled on this Vlan
IGMP snooping immediate-leave is disabled on this Vlan
IGMP snooping mrouter learn mode is pim-dvmrp on this Vlan
IGMP snooping is running in IGMP_ONLY mode on this Vlan
vlan 2
-----
IGMP snooping is globally enabled
IGMP snooping is enabled on this Vlan
IGMP snooping immediate-leave is disabled on this Vlan
IGMP snooping mrouter learn mode is pim-dvmrp on this Vlan
IGMP snooping is running in IGMP_ONLY mode on this Vlan
vlan 10
-----
IGMP snooping is globally enabled
IGMP snooping is enabled on this Vlan
IGMP snooping immediate-leave is disabled on this Vlan
IGMP snooping mrouter learn mode is pim-dvmrp on this Vlan
IGMP snooping is running in IGMP_ONLY mode on this Vlan
```

This is an example of output from the **show ip igmp snooping** privileged EXEC command for a specific VLAN interface:

```
Switch# show ip igmp snooping vlan 1
vlan 1
-----
IGMP snooping is globally enabled
IGMP snooping is disabled on this Vlan
IGMP snooping immediate-leave is disabled on this Vlan
IGMP snooping mrouter learn mode is pim-dvmrp on this Vlan
```

This is an example of output from the **show ip igmp snooping mrouter** privileged EXEC command for VLAN 1:

```
Switch# show ip igmp snooping mrouter vlan 1
Vlan    ports
----    -
1       Gi0/17(dynamic)
1       Gi0/20(dynamic)
```

This example shows how to display the Layer 2 multicast entries for VLAN 1:

```
Switch# show mac address-table multicast vlan 1
vlan  mac address      type      ports
-----+-----+-----+-----+-----
--
1  0100.5e02.0203      user      Gi0/17,Gi0/20
1  0100.5e00.0127      igmp      Gi0/17,Gi0/20
1  0100.5e00.0128      user      Gi0/17,Gi0/20
1  0100.5e00.0001      igmp      Gi0/17,Gi0/20
```

This is an example of output from the **show mac address-table multicast count** privileged EXEC command for the switch:

```
Switch# show mac address-table multicast count
```

```
Multicast MAC Entries for all vlans:    10
```

This is an example of output from the **show mac address-table multicast count** privileged EXEC command for a VLAN:

```
Switch# show mac address-table multicast vlan 1 count
```

```
Multicast MAC Entries for vlan 1:
```

This example shows how to display only the user-configured multicast entries for VLAN 1:

```
Switch# show mac address-table multicast vlan 1 user
```

```
vlan  mac address      type      ports
-----+-----+-----+-----+-----
--
  1  0100.5e02.0203    user      Gi0/17,Gi0/20
  1  0100.5e00.0128    user      Gi0/17,Gi0/20
```

This example shows how to display the total number of entries learned by IGMP snooping for VLAN 1:

```
Switch# show mac address-table multicast vlan 1 igmp-snooping count
```

```
Number of user programmed entries:    2
```

Understanding Multicast VLAN Registration

Multicast VLAN Registration (MVR) is designed for applications using wide-scale deployment of multicast traffic across an Ethernet ring-based service provider network (for example, the broadcast of multiple television channels over a service-provider network). MVR allows a subscriber on a port to subscribe and unsubscribe to a multicast stream on the network-wide multicast VLAN. It allows the single multicast VLAN to be shared in the network while subscribers remain in separate VLANs. MVR provides the ability to continuously send multicast streams in the multicast VLAN, but to isolate the streams from the subscriber VLANs for bandwidth and security reasons.

MVR assumes that subscriber ports subscribe and unsubscribe (join and leave) these multicast streams by sending out IGMP join and leave messages. These messages can originate from an IGMP version-2-compatible host with an Ethernet connection. Although MVR operates on the underlying mechanism of IGMP snooping, the two features operate independently of each other. One can be enabled or disabled without affecting the behavior of the other feature. However, if IGMP snooping and MVR are both enabled, MVR reacts only to join and leave messages from multicast groups configured under MVR. Join and leave messages from all other multicast groups are managed by IGMP snooping.

The switch CPU identifies the MVR IP multicast streams and their associated MAC addresses in the switch forwarding table, intercepts the IGMP messages, and modifies the forwarding table to include or remove the subscriber as a receiver of the multicast

stream, even though the receivers might be in a different VLAN from the source. This forwarding behavior selectively allows traffic to cross between different VLANs.

The switch has these modes of MVR operation: dynamic and compatible.

- When operating in MVR dynamic mode, the switch performs standard IGMP snooping. IGMP information packets are sent to the switch CPU, but multicast data packets are not sent to the CPU. Dynamic mode allows the multicast router to run normally because the switch sends the IGMP join messages to the router, and the router forwards multicast streams for a particular group to an interface only if it has received a join message from the interface for the group. Receiver ports are treated as members of the multicast VLAN for MVR multicast control and data traffic. IGMP reports for MVR groups are sent out source ports in the multicast VLAN.

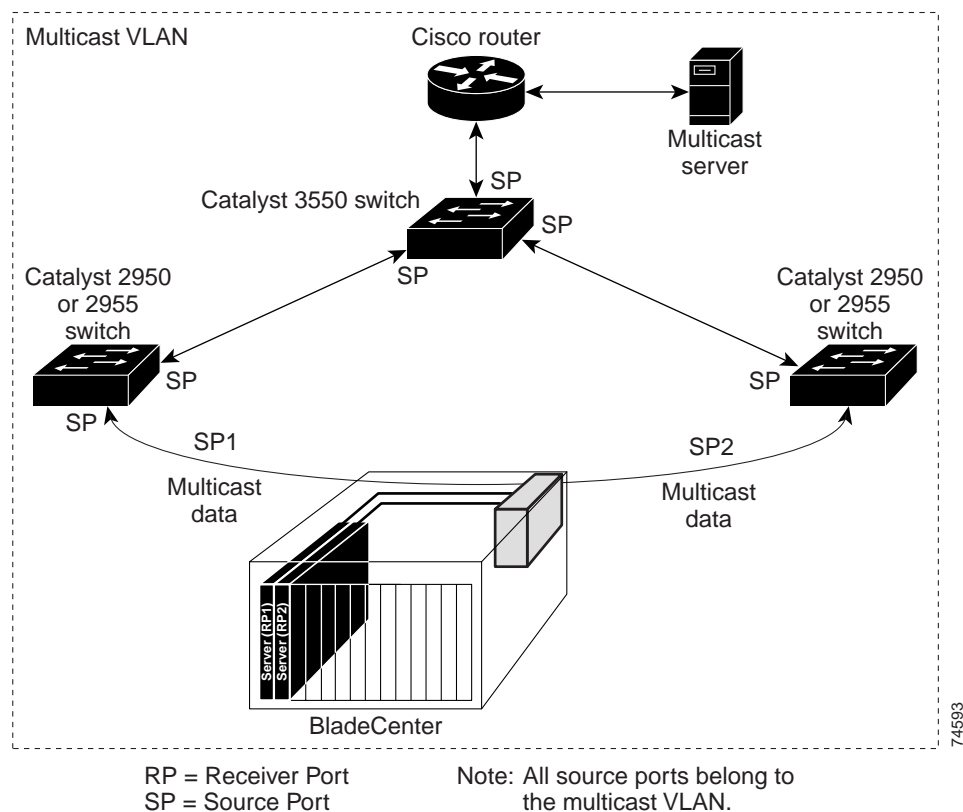
Using MVR in a Multicast Television Application

In a multicast television application, a PC or a television with a set-top box can receive the multicast stream. Multiple set-top boxes or PCs can be connected to one subscriber port, which is a switch port configured as an MVR receiver port. See Figure 58. When a subscriber selects a channel, the set-top box or PC sends an IGMP report to the S1 switch to join the appropriate multicast. If the IGMP report matches one of the configured multicast MAC addresses, the switch CPU modifies the hardware address table to include this receiver port and VLAN as a forwarding destination of the specified multicast stream when it is received from the multicast VLAN. Uplink ports that send and receive multicast data to and from the multicast VLAN are called MVR source ports.

When a subscriber changes channels or turns off the television, the set-top box sends an IGMP leave message for the multicast stream. The switch CPU sends an IGMP group-specific query through the receiver port VLAN. If there is another set-top box in the VLAN still subscribing to this group, that set-top box must respond within the maximum response time. If the CPU does not receive a response, it eliminates the receiver port as a forwarding destination for this group.

If the Immediate-Leave feature is enabled on a receiver port, the port leaves a multicast group more quickly. Without Immediate Leave, when the switch receives an IGMP leave message from a subscriber on a receiver port, it sends out an IGMP query on that port and waits for IGMP group membership reports. If no reports are received in a configured time period, the receiver port is removed from multicast group membership. With Immediate Leave, an IGMP query is not sent from the receiver port on which the IGMP leave was received. As soon as the leave message is received, the receiver port is removed from multicast group membership, which speeds up leave latency. Enable the Immediate Leave feature only on receiver ports to which a single receiver device is connected.

Figure 58. Multicast VLAN Registration Example



MVR eliminates the need to duplicate server multicast traffic for subscribers in each VLAN. Multicast traffic for all servers is only sent around the VLAN trunk once—only on the multicast VLAN. Although the IGMP leave and join message in the VLAN to which the server port is assigned. These messages dynamically register for streams of multicast traffic in the multicast VLAN on the Layer 3 device. The access layer switch (S1 switch) modifies the forwarding behavior to allow the traffic to be forwarded from the multicast VLAN to the server port in a different VLAN, selectively allowing traffic to cross between two VLANs.

IGMP reports are sent to the same MAC addresses as the multicast data. The S1 CPU must capture all IGMP join and leave messages from receiver ports and forward them to the multicast VLAN of the source (uplink) port.

Configuring MVR

These sections include basic MVR configuration information:

- Default MVR Configuration, on page 303
- MVR Configuration Guidelines and Limitations, on page 304
- Configuring MVR Global Parameters, on page 304
- Configuring MVR Interfaces, on page 305

Default MVR Configuration

Table 60 shows the default MVR configuration.

Table 60. Default MVR Configuration

Feature	Default Setting
MVR	Disabled globally and per interface
Multicast addresses	None configured
Query response time	0.5 second
Multicast VLAN	VLAN 1
Mode	Compatible
Interface (per port) default	Neither a receiver nor a source port
Immediate Leave	Disabled on all ports

MVR Configuration Guidelines and Limitations

Follow these guidelines when configuring MVR:

- Receiver ports cannot be trunk ports. Receiver ports on a switch can be in different VLANs, but should not belong to the multicast VLAN.
- The maximum number of multicast entries that can be configured on a switch (that is, the maximum number of television channels that can be received) is 256.
- Each channel is one multicast stream destined for a unique IP multicast address. These IP addresses cannot alias between themselves or with the reserved IP multicast addresses (in the range 224.0.0.xxx).

Note: For complete syntax and usage information for the commands used in this section, refer to the command reference for this release.

Configuring MVR Global Parameters

You do not need to set the optional MVR parameters if you choose to use the default settings. If you do want to change the default parameters (except for the MVR VLAN), you must first enable MVR.

Beginning in privileged EXEC mode, follow these steps to configure MVR parameters:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	mvr	Enable MVR on the switch.
3.	mvr group ip-address [count]	Configure an IP multicast address on the switch or use the <i>count</i> parameter to configure a contiguous series of MVR group addresses (the range for <i>count</i> is 1 to 256; the default is 1). Any multicast data sent to this address is sent to all source ports on the switch and all receiver ports that have elected to receive data on that multicast address. Each multicast address would correspond to one television channel. Note: Each IP address translates to a multicast 48-bit MAC address. If an IP address being configured translates (aliases) to a previously configured MAC address or to any reserved multicast MAC addresses, the command fails.
4.	mvr querytime value	(Optional) Define the maximum time to wait for IGMP report memberships on a receiver port before removing the port from multicast group membership. The value is in units of tenths of a second. The range is from 1 to 100 and the default is 5 tenths or one-half second.

Step	Command	Purpose
5.	mvr vlan <i>vlan-id</i>	(Optional) Specify the VLAN in which multicast data is received; all source ports must belong to this VLAN. The VLAN range is 1 to 4094. The default is VLAN 1.
6.	mvr mode { dynamic compatible }	(Optional) Specify the MVR mode of operation: <ul style="list-style-type: none"> • dynamic—Allows dynamic MVR membership on source ports. • compatible—Is compatible with Catalyst 3500 XL and Catalyst 2900 XL switches and does not support IGMP dynamic joins on source ports. The default is compatible mode.
7.	end	Return to privileged EXEC mode.
8.	show mvr or show mvr members	Verify the configuration.
9.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default settings, use the **no mvr [mode | group ip-address | querytime | vlan]** global configuration commands.

This example shows how to enable MVR, configure the MVR group address, set the query time to 1 second (10 tenths), specify the MVR multicast VLAN as VLAN 22, set the MVR mode as dynamic, and verify the results:

```
Switch(config)# mvr
Switch(config)# mvr group 228.1.23.4
Switch(config)# mvr querytime 10
Switch(config)# mvr vlan 22
Switch(config)# mvr mode dynamic
Switch(config)# end
Switch# show mvr
MVR Running: TRUE
MVR multicast vlan: 22
MVR Max Multicast Groups: 256
MVR Current multicast groups: 1
MVR Global query response time: 10 (tenths of sec)
MVR Mode: dynamic
```

You can use the **show mvr members** privileged EXEC command to verify the MVR multicast group addresses on the switch.

Configuring MVR Interfaces

Beginning in privileged EXEC mode, follow these steps to configure MVR interfaces:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	mvr	Enable MVR on the switch.

Step	Command	Purpose
3.	interface <i>interface-id</i>	Enter interface configuration mode, and enter the type and number of the port to configure; for example, enter gi0/1 or gigabitethernet 0/1 for Gigabit Ethernet port 1.
4.	mvr type { source receiver }	<p>Configure an MVR port as one of these:</p> <ul style="list-style-type: none"> • source—Configure uplink ports that receive and send multicast data as source ports. Subscribers cannot be directly connected to source ports. All source ports on a switch belong to the single multicast VLAN. • receiver—Configure a port as a receiver port if it is a subscriber port and should only receive multicast data. It does not receive data unless it becomes a member of the multicast group, either statically or by using IGMP leave and join messages. Receiver ports cannot belong to the multicast VLAN. <p>The default configuration is as a non-MVR port. If you attempt to configure a non-MVR port with MVR characteristics, the operation fails.</p>
5.	mvr vlan <i>vlan-id</i> group <i>ip-address</i>	<p>(Optional) Statically configure a port to receive multicast traffic sent to the multicast VLAN and the IP multicast address. A port statically configured as a member of a group remains a member of the group until statically removed.</p> <p>Note: In compatible mode, this command applies to only receiver ports. In dynamic mode, it applies to receiver ports and source ports.</p> <p>Receiver ports can also dynamically join multicast groups by using IGMP join and leave messages.</p>
6.	mvr immediate	<p>(Optional) Enable the Immediate Leave feature of MVR on the port.</p> <p>Note: This command applies to only receiver ports and should only be enabled on receiver ports to which a single receiver device is connected.</p>
7.	end	Return to privileged EXEC mode.
8.	show mvr show mvr interface or show mvr members	Verify the configuration.
9.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the interface to its default settings, use the **no mvr [type | immediate | vlan *vlan-id* | group]** interface configuration commands.

This example shows how to configure Gigabit Ethernet port 0/17 as a receiver port, statically configure the port to receive multicast traffic sent to the multicast group address, configure Immediate Leave on the interface, and verify the results.

```
Switch(config)# mvr
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# mvr type receiver
Switch(config-if)# mvr vlan 22 group 228.1.23.4
```

```
Switch(config-if)# mvr immediate
Switch(config)# end
Switch# show mvr interface gigabitethernet0/17
Type: RECEIVER Status: ACTIVE Immediate Leave: ENABLED
```

This is an example of output from the **show mvr interface** privileged EXEC command when the **member** keyword is included:

```
Switch# show mvr interface gigabitethernet0/17 members
224.0.1.1          DYNAMIC ACTIVE
```

Displaying MVR Information

You can display MVR information for the switch or for a specified interface.

Beginning in privileged EXEC mode, use the commands in Table 61 to display MVR configuration:

Table 61. Commands for Displaying MVR Information

show mvr	Displays MVR status and values for the switch—whether MVR is enabled or disabled, the multicast VLAN, the maximum (256) and current (0 through 256) number of multicast groups, the query response time, and the MVR mode.
show mvr interface [<i>interface-id</i>] [members [<i>vlan</i> <i>vlan-id</i>]]	Displays all MVR interfaces and their MVR configurations. When a specific interface is entered, displays this information: <ul style="list-style-type: none"> Type—Receiver or Source Status—One of these: <ul style="list-style-type: none"> Active means the port is part of a VLAN. Up/Down means that the port is forwarding or nonforwarding. Inactive means that the port is not part of any VLAN. Immediate Leave—Enabled or Disabled If the members keyword is entered, displays all multicast group members on this port or, if a VLAN identification is entered, all multicast group members on the VLAN. The VLAN ID range is 1 to 4094.
show mvr members [<i>ip-address</i>]	Displays all receiver and source ports that are members of any IP multicast group or the specified IP multicast group IP address.

This is an example of output from the **show mvr** privileged EXEC command:

```
Switch# show mvr
MVR Running: TRUE
MVR multicast vlan: 1
MVR Max Multicast Groups: 256
MVR Current multicast groups: 256
MVR Global query response time: 5 (tenths of sec)
MVR Mode: compatible
```

This is an example of output from the **show mvr interface** privileged EXEC command:

```
Switch# show mvr interface
Port      Type      Status      Immediate Leave
-----
Gi0/17    SOURCE    ACTIVE/UP    DISABLED
Gi0/18    SOURCE    ACTIVE/UP    DISABLED
Gi0/19    SOURCE    ACTIVE/DOWN  DISABLED
Gi0/20    SOURCE    ACTIVE/DOWN  DISABLED
```

This is an example of output from the **show mvr interface** privileged EXEC command for a specified interface:

```
Switch# show mvr interface gigabitethernet0/20
224.0.1.1      DYNAMIC ACTIVE
```

This is an example of output from the **show mvr interface** privileged EXEC command when the **members** keyword is included:

```
Switch# show mvr interface gigabitethernet0/20 members
224.0.1.1      DYNAMIC ACTIVE
```

This is an example of output from the **show mvr members** privileged EXEC command:

```
Switch# show mvr members
MVR Group IP      Status      Members
-----
224.0.1.1      ACTIVE      Gi0/17(s), Gi0/2(d)
224.0.1.2      ACTIVE      Gi0/17(s)
224.0.1.3      ACTIVE      Gi0/17(s)
224.0.1.4      ACTIVE      Gi0/17(s)
224.0.1.5      ACTIVE      Gi0/17(s)
<output truncated>
```

Configuring IGMP Filtering

In some environments, for example metropolitan or multiple-dwelling unit (MDU) installations, an administrator might want to control the set of multicast groups to which a user on a switch port can belong. This allows the administrator to control the distribution of multicast services, such as IP/TV, based on some type of subscription or service plan. With the IGMP filtering feature, you can filter multicast joins on a per-port basis by configuring IP multicast profiles and associating them with individual switch ports. An IGMP profile can contain one or more multicast groups and specifies whether access to the group is permitted or denied. If an IGMP profile denying access to a multicast group is applied to a switch port, the IGMP join report requesting the stream of IP multicast traffic is dropped, and the port is not allowed to receive IP multicast traffic from that group. If the filtering action permits access to the multicast group, the IGMP report from the port is forwarded for normal processing.

IGMP filtering controls only group specific query and membership reports, including join and leave reports. It does not control general IGMP queries. IGMP filtering has no

relationship with the function that directs the forwarding of IP multicast traffic. The filtering feature operates in the same manner whether CGMP or MVR is used to forward the multicast traffic.

You can also set the maximum number of IGMP groups that a Layer 2 interface can join.

Default IGMP Filtering Configuration

Table 62 shows the default IGMP filtering configuration.

Table 62. Default IGMP Filtering Configuration

Feature	Default Setting
IGMP filters	None applied
IGMP Maximum number of IGMP groups	No maximum set
IGMP profiles	None defined
IGMP profile action	Deny the range addresses

Configuring IGMP Profiles

To configure an IGMP profile, use the **ip igmp profile** global configuration command with a profile number to create an IGMP profile and to enter IGMP profile configuration mode. From this mode, you can specify the parameters of the IGMP profile to be used for filtering IGMP join requests from a port. When you are in IGMP profile configuration mode, you can create the profile by using these commands:

- **deny**
Specifies that matching addresses are denied; this is the default condition.
- **exit**
Exits from igmp-profile configuration mode.
- **no**
Negates a command or sets its defaults.
- **permit**
Specifies that matching addresses are permitted.
- **range**
Specifies a range of IP addresses for the profile. You can enter a single IP address or a range with a start and an end address.

The default is for the switch to have no IGMP profiles configured. When a profile is configured, if neither the **permit** nor **deny** keyword is included, the default is to deny access to the range of IP addresses.

Beginning in privileged EXEC mode, follow these steps to create an IGMP profile:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	ip igmp profile <i>profile number</i>	Enter IGMP profile configuration mode, and assign a number to the profile you are configuring. The range is from 1 to 4294967295.

Step	Command	Purpose
3.	permit deny	(Optional) Set the action to permit or deny access to the IP multicast address. If no action is configured, the default for the profile is to deny access.
4.	range <i>ip multicast address</i>	Enter the IP multicast address or range of IP multicast addresses to which access is being controlled. If entering a range, enter the low IP multicast address, a space, and the high IP multicast address. You can use the range command multiple times to enter multiple addresses or ranges of addresses.
5.	end	Return to privileged EXEC mode.
6.	show ip igmp profile <i>profile number</i>	Verify the profile configuration.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete a profile, use the **no ip igmp profile** *profile number* global configuration command.

To delete an IP multicast address or range of IP multicast addresses, use the **no range** *ip multicast address* IGMP profile configuration command.

This example shows how to create IGMP profile 4 allowing access to the single IP multicast address and how to verify the configuration. If the action was to deny (the default), it would not appear in the **show ip igmp profile** output display.

```
Switch(config)# ip igmp profile 4
Switch(config-igmp-profile)# permit
Switch(config-igmp-profile)# range 229.9.9.0
Switch(config-igmp-profile)# end
Switch# show ip igmp profile 4
IGMP Profile 4
    permit
    range 229.9.9.0 229.9.9.0
```

Applying IGMP Profiles

To control access as defined in an IGMP profile, use the **ip igmp filter** interface configuration command to apply the profile to the appropriate interfaces. You can apply IGMP profiles to Layer 2 ports only. You cannot apply profiles to ports that belong to an EtherChannel port group. You can apply a profile to multiple interfaces, but each interface can only have one profile applied to it.

Beginning in privileged EXEC mode, follow these steps to apply an IGMP profile to a switch port:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and enter the physical interface to configure, for example gigabitethernet0/3 . The interface must be a Layer 2 port that does not belong to an EtherChannel port group.
3.	ip igmp filter <i>profile number</i>	Apply the specified IGMP profile to the interface. The profile number can be from 1 to 4294967295.

Step	Command	Purpose
4.	end	Return to privileged EXEC mode.
5.	show running configuration interface <i>interface-id</i>	Verify the configuration.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a profile from an interface, use the **no ip igmp filter** *profile number* interface configuration command.

This example shows how to apply IGMP profile 4 to an interface and verify the configuration.

```
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# ip igmp filter 4
Switch(config-if)# end
Switch# show running-config interface gigabitethernet0/17
Building configuration...
```

```
Current configuration : 123 bytes
!
interface gigabitethernet0/17
  no ip address
  shutdown
  snmp trap link-status
  ip igmp max-groups 25
  ip igmp filter 4
end
```

Setting the Maximum Number of IGMP Groups

You can set the maximum number of IGMP groups that a Layer 2 interface can join by using the **ip igmp mac-groups** interface configuration command. Use the **no** form of this command to set the maximum back to the default, which is no limit.

You cannot use this command on ports that belong to an EtherChannel port group.

Beginning in privileged EXEC mode, follow these steps to apply an IGMP profile to a switch port:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and enter the physical interface to configure, for example gigabitethernet0/1 . The interface must be a Layer 2 port that does not belong to an EtherChannel group.
3.	ip igmp max-groups <i>number</i>	Set the maximum number of IGMP groups that the interface can join. The range is from 0 to 4294967294. The default is to have no maximum set.
4.	end	Return to privileged EXEC mode.
5.	show running-configuration interface <i>interface-id</i>	Verify the configuration.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the maximum group limitation and return to the default of no maximum, use the **no ip igmp max-groups** interface configuration command.

This example shows how to limit the number of IGMP groups that an interface can join to 25.

```
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# ip igmp max-groups 25
Switch(config-if)# end
Switch# show running-config interface gigabitethernet0/17
Building configuration...
```

```
Current configuration : 123 bytes
!
interface gigabitethernet0/17
 no ip address
 shutdown
 snmp trap link-status
 ip igmp max-groups 25
 ip igmp filter 4
end
```

Displaying IGMP Filtering Configuration

You can display IGMP profile characteristics, and you can display the IGMP profile and maximum group configuration for all interfaces on the switch or for a specified interface.

Use the privileged EXEC commands in Table 63 to display IGMP filtering configuration:

Table 63. Commands for Displaying IGMP Filtering Configuration

show ip igmp profile [<i>profile number</i>]	Displays the specified IGMP profile or all IGMP profiles defined on the switch.
show running-configuration [<i>interface interface-id</i>]	Displays the configuration of the specified interface or all interfaces on the switch, including (if configured) the maximum number of IGMP groups to which an interface can belong and the IGMP profile applied to the interface.

This is an example of the **show ip igmp profile** privileged EXEC command when no profile number is entered. All profiles defined on the switch are displayed.

```
Switch# show ip igmp profile
IGMP Profile 3
  range 230.9.9.0 230.9.9.0
IGMP Profile 4
  permit
  range 229.9.9.0 229.255.255.255
```

This is an example of the output from the **show running-config** privileged EXEC command when an interface is specified with IGMP maximum groups configured and IGMP profile 4 has been applied to the interface.

```
Switch# show running-config interface gigabitethernet0/17
Building configuration...
Current configuration : 123 bytes
!
```

```
interface gigabitethernet0/17
  no ip address
  shutdown
  snmp trap link-status
  ip igmp max-groups 25
  ip igmp filter 4
end
```

Chapter 16. Configuring Port-Based Traffic Control

This chapter describes how to configure the port-based traffic control features on your switch.

Note: For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

This chapter consists of these sections:

- Configuring Storm Control, on page 315
- Configuring Protected Ports, on page 317
- Configuring Port Security, on page 318
- Displaying Port-Based Traffic Control Settings, on page 325

Configuring Storm Control

These sections include storm control configuration information and procedures:

- Understanding Storm Control, on page 315
- Default Storm Control Configuration, on page 316
- Enabling Storm Control, on page 316
- Disabling Storm Control, on page 316

Understanding Storm Control

A packet storm occurs when a large number of broadcast, unicast, or multicast packets are received on a port. Forwarding these packets can cause the network to slow down or to time out. Storm control is configured for the switch as a whole but operates on a per-port basis. By default, storm control is disabled.

Storm control uses rising and falling thresholds to block and then restore the forwarding of broadcast, unicast, or multicast packets. You can also set the switch to shut down the port when the rising threshold is reached.

Storm control uses a bandwidth-based method to measure traffic activity.

The thresholds can either be expressed as a percentage of the total available bandwidth that can be used by the broadcast, multicast, or unicast traffic, or as the rate at which the interface receives multicast, broadcast, or unicast traffic.

When a switch uses the bandwidth-based method, the rising threshold is the percentage of total available bandwidth associated with multicast, broadcast, or unicast traffic before forwarding is blocked. The falling threshold is the percentage of total available bandwidth below which the switch resumes normal forwarding. In general, the higher the level, the less effective the protection against broadcast storms.

When a switch uses traffic rates as the threshold values, the rising and falling thresholds are in packets per second. The rising threshold is the rate at which multicast, broadcast, and unicast traffic is received before forwarding is blocked. The falling threshold is the rate below which the switch resumes normal forwarding. In general, the higher the rate, the less effective the protection against broadcast storms.

Default Storm Control Configuration

By default, broadcast, multicast, and unicast storm control is disabled on the switch. The default action is to filter traffic and to not send an SNMP trap.

Enabling Storm Control

Beginning in privileged EXEC mode, follow these steps to enable storm control:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Specify the port to configure, and enter interface configuration mode.
3.	storm-control { broadcast multicast unicast } level { <i>level</i> [<i>level-low</i>] pps <i>pps</i> <i>pps-low</i> }	Configure broadcast, multicast, or unicast storm control. For <i>level</i> , specify the rising threshold level for broadcast, multicast, or unicast traffic as a percentage of the bandwidth. The storm control action occurs when traffic utilization reaches this level. (Optional) For <i>level-low</i> , specify the falling threshold level as a percentage of the bandwidth. This value must be less than the rising suppression value. The normal transmission restarts (if the action is filtering) when traffic drops below this level. For pps <i>pps</i> , specify the rising threshold level for broadcast, multicast, or unicast traffic in packets per second. The storm control action occurs when traffic reaches this level. For <i>pps</i> and <i>pps-low</i> , the range is from 0 to 4294967295.
4.	storm-control action { shutdown trap }	Specify the action to be taken when a storm is detected. The default is to filter out the traffic and not to send traps. Select the shutdown keyword to error-disable the port during a storm. Select the trap keyword to generate an SNMP trap when a storm is detected.
5.	end	Return to privileged EXEC mode.
6.	show storm-control [<i>interface</i>] [{ broadcast history multicast unicast }]	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The output from the **show storm-control** privileged EXEC command shows the upper, lower, and current thresholds as a percentage of the total bandwidth or the packets per second, depending on the configuration.

Disabling Storm Control

Beginning in privileged EXEC mode, follow these steps to disable storm control:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Specify the port to configure, and enter interface configuration mode.
3.	no storm-control {broadcast multicast unicast} level	Disable port storm control.
4.	no storm-control action {shutdown trap}	Disable the specified storm control action.
5.	end	Return to privileged EXEC mode.
6.	show storm-control {broadcast multicast unicast}	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Configuring Protected Ports

Some applications require that no traffic be forwarded between ports on the same switch so that one neighbor does not see the traffic generated by another neighbor. In such an environment, the use of protected ports ensures that there is no exchange of unicast, broadcast, or multicast traffic between these ports on the switch.

Protected ports have these features:

- A protected port does not forward any traffic (unicast, multicast, or broadcast) to any other port that is also a protected port. Traffic cannot be forwarded between protected ports at Layer 2; all traffic passing between protected ports must be forwarded through a Layer 3 device.
- Forwarding behavior between a protected port and a nonprotected port proceeds as usual.
- Protected ports are supported on 802.1Q trunks.

The default is to have no protected ports defined.

You can configure protected ports on a physical interface (for example, Gigabit Ethernet 0/17) or an EtherChannel group (for example, port-channel 5). When you enable protected ports for a port channel, it is enabled for all ports in the port-channel group.

Beginning in privileged EXEC mode, follow these steps to define a port as a protected port:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Specify the type and number of the physical interface to configure, for example gigabitethernet0/17 , and enter interface configuration mode.
3.	switchport protected	Configure the interface to be a protected port.
4.	end	Return to privileged EXEC mode.
5.	show interfaces <i>interface-id</i> switchport	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable protected port, use the **no switchport protected** interface configuration command.

This example shows how to configure Gigabit Ethernet interface 0/17 as a protected port and verify the configuration:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# switchport protected
Switch(config-if)# end
Switch# show interfaces gigabitethernet0/17 switchport
Name: Gi0/17
Switchport: Enabled

<output truncated>

Protected: True
Unknown unicast blocked: disabled
Unknown multicast blocked: disabled
```

Configuring Port Security

You can use the port security feature to restrict input to an interface by limiting and identifying MAC addresses of the stations allowed to access the port. When you assign secure MAC addresses to a secure port, the port does not forward packets with source addresses outside the group of defined addresses.

Note: You cannot configure port security on the internal 100 Mbps management module ports.

This section includes information about these topics:

- Understanding Port Security, on page 318
- Default Port Security Configuration, on page 320
- Port Security Configuration Guidelines, on page 320
- Enabling and Configuring Port Security, on page 321
- Enabling and Configuring Port Security Aging, on page 323

Understanding Port Security

This section includes information about:

- Secure MAC Addresses, on page 318
- Security Violations, on page 319

Secure MAC Addresses

You can configure these types of secure MAC addresses:

- Static secure MAC addresses—These are manually configured by using the **switchport port-security mac-address *mac-address*** interface configuration command, stored in the address table, and added to the switch running configuration.
- Dynamic secure MAC addresses—These are dynamically learned, stored only in the address table, and removed when the switch restarts.
- *Sticky* secure MAC addresses—These can be dynamically learned or manually configured, stored in the address table, and added to the running configuration. If these addresses are saved in the configuration file, the interface does not need to dynamically relearn them when the switch restarts. Although sticky secure addresses can be manually configured, we do not recommend it.

You can configure an interface to convert the dynamic MAC addresses to sticky secure MAC addresses and to add them to the running configuration by enabling *sticky learning*. To enable sticky learning, enter the **switchport port-security mac-address sticky** interface configuration command. When you enter this command, the interface converts all the dynamic secure MAC addresses, including those that were dynamically learned before sticky learning was enabled, to sticky secure MAC addresses.

The sticky secure MAC addresses do not automatically become part of the configuration file, which is the startup configuration used each time the switch restarts. If you save the sticky secure MAC addresses in the configuration file, when the switch restarts, the interface does not need to relearn these addresses. If you do not save the configuration, they are lost.

If sticky learning is disabled, the sticky secure MAC addresses are converted to dynamic secure addresses and are removed from the running configuration.

A secure port can have from 1 to 132 associated secure addresses. The total number of available secure addresses on the switch is 1024.

Security Violations

It is a security violation when one of these situations occurs:

- The maximum number of secure MAC addresses have been added to the address table, and a station whose MAC address is not in the address table attempts to access the interface.
- An address learned or configured on one secure interface is seen on another secure interface in the same VLAN.

You can configure the interface for one of three violation modes, based on the action to be taken if a violation occurs:

- **protect**—When the number of secure MAC addresses reaches the limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses or increase the number of maximum allowable addresses. You are not notified that a security violation has occurred.
- **restrict**—When the number of secure MAC addresses reaches the limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses or increase the number of maximum allowable addresses. In this mode, you are notified that a security violation has occurred. Specifically, an SNMP trap is sent, a syslog message is logged, and the violation counter increments.
- **shutdown**—In this mode, a port security violation causes the interface to immediately become error-disabled, and turns off the port LED. It also sends an SNMP trap, logs a syslog message, and increments the violation counter. When a secure port is in the error-disabled state, you can bring it out of this state by entering the **errdisable recovery cause psecure-violation** global configuration command, or you can manually re-enable it by entering the **shutdown** and **no shutdown** interface configuration commands. This is the default mode.

Table 64 shows the violation mode and the actions taken when you configure an interface for port security.

Table 64. Security Violation Mode Actions

Violation Mode	Traffic is forwarded ¹	Sends SNMP trap	Sends syslog message	Displays error message ²	Violation counter increments	Shuts down port
protect	No	No	No	No	No	No
restrict	No	Yes	Yes	No	Yes	No
shutdown	No	Yes	Yes	No	Yes	Yes

1. Packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses.

2. The switch will return an error message if you manually configure an address that would cause a security violation.

Default Port Security Configuration

Table 65 shows the default port security configuration for an interface.

Table 65. Default Port Security Configuration

Feature	Default Setting
Port security	Disabled.
Maximum number of secure MAC addresses	One.
Violation mode	Shutdown.
Sticky address learning	Disabled.
Port security aging	Disabled. Aging time is 0. When enabled, the default type is absolute .

Port Security Configuration Guidelines

Follow these guidelines when configuring port security:

- Port security can only be configured on static access ports.
- A secure port cannot be a dynamic access port or a trunk port.
- A secure port cannot be a destination port for Switched Port Analyzer (SPAN).
- A secure port cannot belong to a Fast EtherChannel or Gigabit EtherChannel port group.
- You cannot configure static secure or sticky secure MAC addresses on a voice VLAN.
- When you enable port security on an interface that is also configured with a voice VLAN, you must set the maximum allowed secure addresses on the port to at least two.
- If any type of port security is enabled on the access VLAN, dynamic port security is automatically enabled on the voice VLAN.
- When a voice VLAN is configured on a secure port that is also configured as a sticky secure port, all addresses seen on the voice VLAN are learned as dynamic secure addresses, and all addresses seen on the access VLAN (to which the port belongs) are learned as sticky secure addresses.
- You cannot configure port security on a per-VLAN basis.
- The switch does not support port security aging of sticky secure MAC addresses.
- The **protect** and **restrict** options cannot be simultaneously enabled on an interface.

Enabling and Configuring Port Security

Beginning in privileged EXEC mode, follow these steps to restrict input to an interface by limiting and identifying MAC addresses of the stations allowed to access the port:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Specify the type and number of the physical interface to configure, for example gigabitethernet0/17 , and enter interface configuration mode.
3.	switchport mode access	Set the interface mode as access ; an interface in the default mode (dynamic desirable) cannot be configured as a secure port.
4.	switchport port-security	Enable port security on the interface.
5.	switchport port-security maximum <i>value</i>	(Optional) Set the maximum number of secure MAC addresses for the interface. The range is 1 to 132; the default is 1.
6.	switchport port-security violation { protect restrict shutdown }	<p>(Optional) Set the violation mode, the action to be taken when a security violation is detected, as one of these:</p> <ul style="list-style-type: none"> • protect—When the number of secure MAC addresses reaches the limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses or increase the number of maximum allowable addresses. You are not notified that a security violation has occurred. • restrict—When the number of secure MAC addresses reaches the limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses or increase the number of maximum allowable addresses. In this mode, you are notified that a security violation has occurred. Specifically, an SNMP trap is sent, a syslog message is logged, and the violation counter increments. • shutdown—In this mode, a port security violation causes the interface to immediately become error-disabled, and turns off the port LED. It also sends an SNMP trap, logs a syslog message, and increments the violation counter. <p>Note: When a secure port is in the error-disabled state, you can bring it out of this state by entering the errdisable recovery cause psecure-violation global configuration command, or you can manually re-enable it by entering the shutdown and no shutdown interface configuration commands.</p>

Step	Command	Purpose
7.	switchport port-security mac-address <i>mac-address</i>	(Optional) Enter a static secure MAC address for the interface, repeating the command as many times as necessary. You can use this command to enter the maximum number of secure MAC addresses. If you configure fewer secure MAC addresses than the maximum, the remaining MAC addresses are dynamically learned. Note: If you enable sticky learning after you enter this command, the secure addresses that were dynamically learned are converted to sticky secure MAC addresses and are added to the running configuration.
8.	switchport port-security mac-address sticky	(Optional) Enable sticky learning on the interface.
9.	end	Return to privileged EXEC mode.
10.	show port-security	Verify your entries.
11.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the interface to the default condition as not a secure port, use the **no switchport port-security** interface configuration command. If you enter this command when sticky learning is enabled, the sticky secure addresses remain part of the running configuration but are removed from the address table. All addresses are now dynamically learned.

To return the interface to the default number of secure MAC addresses, use the **no switchport port-security maximum** *value* interface configuration command.

To return the violation mode to the default condition (shutdown mode), use the **no switchport port-security violation {protect | restrict}** interface configuration command.

To disable sticky learning on an interface, use the **no switchport port-security mac-address sticky** interface configuration command. The interface converts the sticky secure MAC addresses to dynamic secure addresses.

To delete a static secure MAC address from the address table, use the **clear port-security configured address** *mac-address* privileged EXEC command. To delete all the static secure MAC addresses on an interface, use the **clear port-security configured interface** *interface-id* privileged EXEC command.

To delete a dynamic secure MAC address from the address table, use the **clear port-security dynamic address** *mac-address* privileged EXEC command. To delete all the dynamic addresses on an interface, use the **clear port-security dynamic interface** *interface-id* privileged EXEC command.

To delete a sticky secure MAC addresses from the address table, use the **clear port-security sticky address** *mac-address* privileged EXEC command. To delete all the sticky addresses on an interface, use the **clear port-security sticky interface** *interface-id* privileged EXEC command.

This example shows how to enable port security on Gigabit Ethernet port 17 and to set the maximum number of secure addresses to 50. The violation mode is the default, no static secure MAC addresses are configured, and sticky learning is enabled.

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet0/17
```

```

Switch(config-if)# switchport mode access
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security maximum 50
Switch(config-if)# switchport port-security mac-address sticky
Switch(config-if)# end
Switch# show port-security interface gigabitethernet0/17
Port Security           : Enabled
Port Status             : Secure-up
Violation Mode          : Shutdown
Aging Time               : 20 mins
Aging Type               : Inactivity
SecureStatic Address Aging : Enabled
Maximum MAC Addresses   : 50
Total MAC Addresses     : 11
Configured MAC Addresses : 0
Sticky MAC Addresses    : 11
Last Source Address     : 0000.0000.0000
Security Violation Count : 0

```

This example shows how to configure a static secure MAC address on Gigabit Ethernet port 17, enable sticky learning, and verify the configuration:

```

Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# switchport mode access
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security mac-address 0000.02000.0004
Switch(config-if)# switchport port-security mac-address sticky
Switch(config-if)# end
Switch# show port-security address
=          Secure Mac Address Table

```

Vlan	Mac Address	Type	Ports	Remaining Age (mins)
1	0000.0000.000a	SecureDynamic	Gi0/17	-
1	0000.0002.0300	SecureDynamic	Gi0/17	-
1	0000.0200.0003	SecureConfigured	Gi0/17	-
1	0000.0200.0004	SecureConfigured	Gi0/18	-
1	0003.fd62.1d40	SecureConfigured	Gi0/19	-
1	0003.fd62.1d45	SecureConfigured	Gi0/19	-
1	0003.fd62.21d3	SecureSticky	Gi0/19	-
1	0005.7428.1a45	SecureSticky	Gi0/20	-
1	0005.7428.1a46	SecureSticky	Gi0/20	-
1	0006.1218.2436	SecureSticky	Gi0/20	-

```

Total Addresses in System :10
Max Addresses limit in System :1024

```

Enabling and Configuring Port Security Aging

You can use port security aging to set the aging time for static and dynamic secure addresses on a port. Two types of aging are supported per port:

- **Absolute**—The secure addresses on the port are deleted after the specified aging time.

- Inactivity—The secure addresses on the port are deleted only if the secure addresses are inactive for the specified aging time.

Use this feature to remove and add PCs on a secure port without manually deleting the existing secure MAC addresses and to still limit the number of secure addresses on a port. You can enable or disable the aging of statically-configured secure addresses on a per-port basis.

Beginning in privileged EXEC mode, follow these steps to configure port security aging:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Specify the port on which you want to enable port security aging, and enter interface configuration mode. Note: The switch does not support port security aging of sticky secure addresses.
3.	switchport port-security aging { static time <i>time</i> type { absolute inactivity }}	Enable or disable static aging for the secure port, or set the aging time or type. Enter static to enable aging for statically configured secure addresses on this port. For <i>time</i> , specify the aging time for this port. The valid range is from 0 to 1440 minutes. If the time is equal to 0, aging is disabled for this port. For type , select one of these keywords: <ul style="list-style-type: none"> • absolute—Sets the aging type as absolute aging. All the secure addresses on this port age out after the specified time (minutes) lapses and are removed from the secure address list. Note: The absolute aging time could vary by 1 minute, depending on the sequence of the system timer. <ul style="list-style-type: none"> • inactivity—Sets the aging type as inactivity aging. The secure addresses on this port age out only if there is no data traffic from the secure source addresses for the specified time period.
4.	end	Return to privileged EXEC mode.
5.	show port-security [interface <i>interface-id</i>] [address]	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable port security aging for all secure addresses on a port, use the **no switchport port-security aging time** interface configuration command. To disable aging for only statically configured secure addresses, use the **no switchport port-security aging static** interface configuration command.

This example shows how to set the aging time as 2 hours for the secure addresses on the Gigabit Ethernet interface 0/17:

```
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# switchport port-security aging time 120
```

This example shows how to set the aging time as 2 minutes for the inactivity aging type with aging enabled for the configured secure addresses on the interface:

```
Switch(config-if)# switchport port-security aging time 2
Switch(config-if)# switchport port-security aging type inactivity
Switch(config-if)# switchport port-security aging static
```

You can verify the previous commands by entering the **show port-security interface interface-id** privileged EXEC command.

Displaying Port-Based Traffic Control Settings

The **show interfaces interface-id switchport** privileged EXEC command displays (among other characteristics) the interface traffic suppression and control configuration. The **show interfaces counters** privileged EXEC commands display the count of discarded packets. The **show storm-control** and **show port-security** privileged EXEC commands display those features.

To display traffic control information, use one or more of the privileged EXEC commands in Table 66.

Table 66. Commands for Displaying Traffic Control Status and Configuration

Command	Purpose
show interfaces [interface-id] switchport	Displays the administrative and operational status of all switching (nonrouting) ports or the specified port, including port protection settings.
show storm-control [interface-id] [broadcast multicast unicast]	Displays storm control suppression levels set on all interfaces or the specified interface for the specified traffic type or for broadcast traffic if no traffic type is entered.
show interfaces [interface-id] counters broadcast	Displays the storm-control broadcast suppression discard counter with the number of packets discarded for all interfaces or the specified interface.
show interfaces [interface-id] counters multicast	Displays the storm-control multicast suppression discard counter with the number of packets discarded for all interfaces or the specified interface.
show interfaces [interface-id] counters unicast	Displays the storm-control unicast suppression discard counter with the number of packets discarded for all interfaces or the specified interface.
show port-security [interface interface-id]	Displays port security settings for the switch or for the specified interface, including the maximum allowed number of secure MAC addresses for each interface, the number of secure MAC addresses on the interface, the number of security violations that have occurred, and the violation mode.
show port-security [interface interface-id] address	Displays all secure MAC addresses configured on all switch interfaces or on a specified interface with aging information for each address.

Chapter 17. Configuring UDLD

This chapter describes how to configure the UniDirectional Link Detection (UDLD) protocol on your switch.

Note: For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference*.

This chapter consists of these sections:

- Understanding UDLD, on page 327
- Configuring UDLD, on page 328
- Displaying UDLD Status, on page 330

Understanding UDLD

UDLD is a Layer 2 protocol that enables devices connected through twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists.

This feature is supported on the external 10/100/1000 Mbps switch ports only. All connected devices must support UDLD for the protocol to successfully identify and disable unidirectional links. When UDLD detects a unidirectional link, it administratively shuts down the affected port and alerts you. Unidirectional links can cause a variety of problems, including spanning-tree topology loops.

UDLD works with the Layer 1 mechanisms to determine the physical status of a link. At Layer 1, autonegotiation takes care of physical signaling and fault detection. UDLD performs tasks that autonegotiation cannot perform, such as detecting the identities of neighbors and shutting down misconnected interfaces. When you enable both autonegotiation and UDLD, Layer 1 and Layer 2 detections work together to prevent physical and logical unidirectional connections and the malfunctioning of other protocols.

A unidirectional link occurs whenever traffic sent by the local device is received by the neighbor but traffic from the neighbor is not received by the local device. If one of the strands in a pair is disconnected, as long as autonegotiation is active, the link does not stay up. In this case, the logical link is undetermined, and UDLD does not take any action. If both strands are working normally from a Layer 1 perspective, UDLD at Layer 2 determines whether those strands are connected correctly and whether traffic is flowing bidirectionally between the correct neighbors. This check cannot be performed by autonegotiation because autonegotiation operates at Layer 1.

UDLD operates by using two mechanisms:

- Neighbor database maintenance

UDLD learns about other UDLD-capable neighbors by periodically sending a hello packet (also called an advertisement or probe) on every active interface to keep each device informed about its neighbors.

When the switch receives a hello message, it caches the information until the age time (hold time or time-to-live) expires. If the switch receives a new hello message before an older cache entry ages, the switch replaces the older entry with the new one.

Whenever an interface is disabled and UDLD is running, whenever UDLD is disabled on an interface, or whenever the switch is reset, UDLD clears all existing cache entries for the interfaces affected by the configuration change. UDLD sends at least one message to inform the neighbors to flush the part of their caches


affected by the status change. The message is intended to keep the caches synchronized.

- Event-driven detection and echoing

UDLD relies on echoing as its detection mechanism. Whenever a UDLD device learns about a new neighbor or receives a resynchronization request from an out-of-sync neighbor, it restarts the detection window on its side of the connection and sends echo messages in reply. Because this behavior is the same on all UDLD neighbors, the sender of the echoes expects to receive an echo in reply. If the detection window ends and no valid reply message is received, the link is considered unidirectional, and the interface is shut down.

Figure 59 shows an example of a unidirectional link condition.

Figure 59. UDLD Detection of a Unidirectional Link



Configuring UDLD

This section describes how to configure UDLD on your switch. It contains this configuration information:

- Default UDLD Configuration, on page 328
- Enabling UDLD Globally, on page 329
- Enabling UDLD on an Interface, on page 329
- Resetting an Interface Shut Down by UDLD, on page 330

Default UDLD Configuration

Table 67 shows the default UDLD configuration.

Table 67. Default UDLD Configuration

Feature	Default Setting
UDLD global enable state	Globally disabled
UDLD per-interface enable state for twisted-pair (copper) media	Disabled on all Ethernet 100 Mbps and 1000BASE-TX interfaces
UDLD aggressive mode	Disabled

A UDLD-capable interface also cannot detect a unidirectional link if it is connected to a UDLD-incapable port of another switch.

Enabling UDLD Globally

Note: This feature and its commands are not supported on the switch. To use the UDLD feature on the switch, see the “Enabling UDLD on an Interface” section on page 329.

Beginning in privileged EXEC mode, follow these steps to enable UDLD in the aggressive or normal mode and to set the configurable message timer on all fiber-optic interfaces on the switch:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	udld { aggressive enable message time <i>message-timer-interval</i> }	Specify the UDLD mode of operation: <ul style="list-style-type: none"> • aggressive—Enables UDLD in aggressive mode on all fiber-optic interfaces. For details on the usage guidelines for the aggressive mode, refer to the command reference guide. • enable—Enables UDLD in normal mode on all fiber-optic interfaces on the switch. UDLD is disabled by default. An individual interface configuration overrides the setting of the udld enable global configuration command. • message time <i>message-timer-interval</i>—Configures the period of time between UDLD probe messages on ports that are in the advertisement phase and are determined to be bidirectional. The range is from 7 to 90 seconds. <p>Note: This command affects fiber-optic interfaces only. Use the udld interface configuration command to enable UDLD on other interface types. For more information, see the “Enabling UDLD on an Interface” section on page 329.</p>
3.	end	Return to privileged EXEC mode.
4.	show udld	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable UDLD globally, use the **no udld enable** global configuration command to disable normal mode UDLD on all fiber-optic ports. Use the **no udld aggressive** global configuration command to disable aggressive mode UDLD on all fiber-optic ports.

Enabling UDLD on an Interface

Beginning in privileged EXEC mode, follow these steps to enable UDLD in the aggressive or normal mode on an interface:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be enabled for UDLD.

Step	Command	Purpose
3.	udld {aggressive enable}	Specify the UDLD mode of operation: <ul style="list-style-type: none"> • aggressive—Enables UDLD in aggressive mode on the specified interface. For details on the usage guidelines for the aggressive mode, refer to the command reference for this release. • enable—Enables UDLD in normal mode on the specified interface. UDLD is disabled by default.
4.	end	Return to privileged EXEC mode.
5.	show udld <i>interface-id</i>	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable UDLD on an interface, use the **no udld enable** interface configuration command.

Resetting an Interface Shut Down by UDLD

Beginning in privileged EXEC mode, follow these steps to reset all interfaces shut down by UDLD:

Step	Command	Purpose
1.	udld reset	Reset all interfaces shut down by UDLD.
2.	show udld	Verify your entries.
3.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

You can also bring up the interface by using these commands:

- The **shutdown** interface configuration command followed by the **no shutdown** interface configuration command restarts the disabled interface.
- The **udld disable** interface configuration command followed by the **udld {aggressive | enable}** interface configuration command re-enables UDLD on the specified interface.
- The **errdisable recovery cause udld** global configuration command enables the timer to automatically recover from the UDLD error-disabled state, and the **errdisable recovery interval *interval*** global configuration command specifies the time to recover from the UDLD error-disabled state.

Displaying UDLD Status

To display the UDLD status for the specified interface or for all interfaces, use the **show udld [*interface-id*]** privileged EXEC command.

For detailed information about the fields in the display, refer to the command reference for this release.

Chapter 18. Configuring CDP

This chapter describes how to configure Cisco Discovery Protocol (CDP) on your switch.

Note: For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference*.

This chapter consists of these sections:

- Understanding CDP, on page 331
- Configuring CDP, on page 331
- Monitoring and Maintaining CDP, on page 334

Understanding CDP

CDP is a device discovery protocol that runs over Layer 2 (the data link layer) on all Cisco-manufactured devices (routers, bridges, access servers, and switches) and allows network management applications to discover Cisco devices that are neighbors of already known devices. With CDP, network management applications can learn the device type and the Simple Network Management Protocol (SNMP) agent address of neighboring devices running lower-layer, transparent protocols. This feature enables applications to send SNMP queries to neighboring devices.

CDP runs on all media that support Subnetwork Access Protocol (SNAP). Because CDP runs over the data-link layer only, two systems that support different network-layer protocols can learn about each other.

Each CDP-configured device sends periodic messages to a multicast address, advertising at least one address at which it can receive SNMP messages. The advertisements also contain time-to-live, or holdtime information, which is the length of time a receiving device holds CDP information before discarding it. Each device also listens to the messages sent by other devices to learn about neighboring devices.

On the switch, CDP enables the Cluster Management Suite to display a graphical view of the network. The switch uses CDP to find cluster candidates and maintain information about cluster members and other devices up to three cluster-enabled devices away from the command switch by default.

The switch supports CDP version 2.

Configuring CDP

These sections include CDP configuration information and procedures:

- Default CDP Configuration, on page 331
- Configuring the CDP Characteristics, on page 332
- Disabling and Enabling CDP, on page 333
- Disabling and Enabling CDP on an Interface, on page 333

Default CDP Configuration

Table 68 shows the default CDP configuration.

Table 68. Default CDP Configuration

Feature	Default Setting
CDP global state	Enabled
CDP interface state	Enabled
CDP timer (packet update frequency)	60 seconds
CDP holdtime (before discarding)	180 seconds
CDP version-2 advertisements	Enabled

Configuring the CDP Characteristics

You can configure the frequency of CDP updates, the amount of time to hold the information before discarding it, and whether or not to send version-2 advertisements.

Beginning in privileged EXEC mode, follow these steps to configure the CDP timer, holdtime, and advertisement type.

Note: Steps 2 through 4 are all optional and can be performed in any order.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	cdp timer <i>seconds</i>	(Optional) Set the transmission frequency of CDP updates in seconds. The range is from 5 to 254; the default is 60 seconds.
3.	cdp holdtime <i>seconds</i>	(Optional) Specify the amount of time a receiving device should hold the information sent by your device before discarding it. The range is from 10 to 255 seconds; the default is 180 seconds.
4.	cdp advertise-v2	(Optional) Configure CDP to send version-2 advertisements. This is the default state.
5.	end	Return to privileged EXEC mode.
6.	show cdp	Verify configuration by displaying global information about CDP on the device.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no** form of the CDP commands to return to the default settings.

This example shows how to configure and verify CDP characteristics.

```
Switch# configure terminal
Switch(config)# cdp timer 50
Switch(config)# cdp holdtime 120
Switch(config)# cdp advertise-v2
Switch(config)# end
```

```
Switch# show cdp
```

```
Global CDP information:
```

```
    Sending CDP packets every 50 seconds
```

```
    Sending a holdtime value of 120 seconds
```

For additional CDP **show** commands, see the “Monitoring and Maintaining CDP” section on page 334.

Disabling and Enabling CDP

CDP is enabled by default.

Note: Creating and maintaining switch clusters is based on the regular exchange of CDP messages. Disabling CDP can interrupt cluster discovery. For more information, see Chapter 5 “Clustering Switches.”

Beginning in privileged EXEC mode, follow these steps to disable the CDP device discovery capability:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	no cdp run	Disable CDP.
3.	end	Return to privileged EXEC mode.

Beginning in privileged EXEC mode, follow these steps to enable CDP when it has been disabled:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	cdp run	Enable CDP after disabling it.
3.	end	Return to privileged EXEC mode.

This example shows how to enable CDP if it has been disabled.

```
Switch# configure terminal
Switch(config)# cdp run
Switch(config)# end
```

Disabling and Enabling CDP on an Interface

CDP is enabled by default on all supported interfaces to send and receive CDP information.

Beginning in privileged EXEC mode, follow these steps to disable CDP on an interface:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and enter the interface on which you are disabling CDP.
3.	no cdp enable	Disable CDP on an interface.
4.	end	Return to privileged EXEC mode.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Beginning in privileged EXEC mode, follow these steps to enable CDP on an interface when it has been disabled:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and enter the interface on which you are enabling CDP.
3.	cdp enable	Enable CDP on an interface after disabling it.
4.	end	Return to privileged EXEC mode.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to enable CDP on an interface when it has been disabled.

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# cdp enable
Switch(config-if)# end
```

Monitoring and Maintaining CDP

To monitor and maintain CDP on your device, perform one or more of these tasks, beginning in privileged EXEC mode.

Command	Description
clear cdp counters	Reset the traffic counters to zero.
clear cdp table	Delete the CDP table of information about neighbors.
show cdp	Display global information, such as frequency of transmissions and the holdtime for packets being sent.
show cdp entry <i>entry-name</i> [protocol version]	Display information about a specific neighbor. You can enter an asterisk (*) to display all CDP neighbors, or you can enter the name of the neighbor about which you want information. You can also limit the display to information about the protocols enabled on the specified neighbor or information about the version of software running on the device.
show cdp interface [<i>type number</i>]	Display information about interfaces where CDP is enabled. You can limit the display to the type of interface or the number of the interface about which you want information (for example, entering gigabitethernet 0/17 displays information only about Gigabit Ethernet port 1).
show cdp neighbors [<i>type number</i>] [detail]	Display information about neighbors, including device type, interface type and number, holdtime settings, capabilities, platform, and port ID. You can limit the display to neighbors on a specific type or number of interface or expand the display to provide more detailed information.
show cdp traffic	Display CDP counters, including the number of packets sent and received and checksum errors.

This is an example of the output from the **show cdp** privileged EXEC commands:

```
Switch# show cdp
```


Global CDP information:

Sending CDP packets every 50 seconds
Sending a holdtime value of 120 seconds
Sending CDPv2 advertisements is enabled

Chapter 19. Configuring SPAN and RSPAN

This chapter describes how to configure Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) on your switch.

Note: For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

This chapter consists of these sections:


- Understanding SPAN and RSPAN, on page 337
- Configuring SPAN, on page 342
- Configuring RSPAN, on page 347
- Displaying SPAN and RSPAN Status, on page 351

Understanding SPAN and RSPAN

You can analyze network traffic passing through ports by using SPAN to send a copy of the traffic to another port on the switch that has been connected to a SwitchProbe device or other Remote Monitoring (RMON) probe or security device. SPAN mirrors received or transmitted (or both) traffic on one or more source ports to a destination port for analysis.

For example, in Figure 60, all traffic on port 5 (the source port) is mirrored to port 17 (the destination port). A network analyzer on port 17 receives all network traffic from port 5 without being physically attached to port 5.


Figure 60. Example SPAN Configuration



Only traffic that enters or leaves source ports can be monitored by using SPAN.

RSPAN extends SPAN by enabling remote monitoring of multiple switches across your network. The traffic for each RSPAN session is carried over a user-specified RSPAN VLAN that is dedicated for that RSPAN session in all participating switches. The SPAN traffic from the sources is copied onto the RSPAN VLAN through a reflector port and then forwarded over trunk ports that are carrying the RSPAN VLAN to any RSPAN destination session monitoring the RSPAN VLAN, as shown in Figure 61.

Figure 61. Example of RSPAN Configuration



SPAN and RSPAN do not affect the switching of network traffic on source ports; a copy of the packets received or sent by the source interfaces are sent to the destination interface. Except for traffic that is required for the SPAN or RSPAN session, reflector ports and destination ports do not receive or forward traffic.

You can use the SPAN destination port to inject traffic from a network security device. For example, if you connect a Cisco Intrusion Detection System (IDS) Sensor Appliance to a destination port, the IDS device can send TCP Reset packets to close down the TCP session of a suspected attacker.

Note: You cannot use the RSPAN destination port to inject traffic from a network security device. The switch does not support ingress forwarding on an RSPAN destination port.

SPAN and RSPAN Concepts and Terminology

This section describes concepts and terminology associated with SPAN and RSPAN configuration.

SPAN Session

A local SPAN session is an association of a destination port with source ports. You can monitor incoming or outgoing traffic on a series or range of ports.

An RSPAN session is an association of source ports across your network with an RSPAN VLAN. The destination source is the RSPAN VLAN.

SPAN sessions do not interfere with the normal operation of the switch. However, an oversubscribed SPAN destination, for example, a 10-Mbps port monitoring a 100-Mbps port, results in dropped or lost packets.

You can configure SPAN sessions on disabled ports; however, a SPAN session does not become active unless you enable the destination port and at least one source port for that session. The **show monitor session session_number** privileged EXEC command displays the operational status of a SPAN session.

A SPAN session remains inactive after system power-on until the destination port is operational.

Traffic Types

SPAN sessions include these traffic types:

- Receive (Rx) SPAN—The goal of receive (or ingress) SPAN is to monitor as much as possible all the packets received by the source interface. A copy of each packet received by the source is sent to the destination port for that SPAN session. You can monitor a series or range of ingress ports in a SPAN session.

At the destination port, if tagging is enabled, the packets appear with the 802.1Q header. If no tagging is specified, packets appear in the native format.

Packets that are modified because of quality of service (QoS)—for example, modified Differentiated Services Code Point (DSCP)—are copied with modification for Rx SPAN.

- **Transmit (Tx) SPAN**—The goal of transmit (or egress) SPAN is to monitor as much as possible all the packets sent by the source interface after all modification and processing is performed by the switch. A copy of each packet sent by the source is sent to the destination port for that SPAN session. The copy is provided after the packet is modified. You can monitor a range of egress ports in a SPAN session.

Packets that are modified because of QoS, the modified packet might not have the same DSCP (IP packet) or CoS (non-IP packet) as the SPAN source.

Some features that can cause a packet to be dropped during transmit processing might also affect the duplicated copy for SPAN. These features include IP standard and extended output ACLs on multicast packets, and egress QoS policing. In the case of output ACLs, if the SPAN source drops the packet, the SPAN destination would also drop the packet. If the source port is oversubscribed, the destination ports will have different dropping behavior.

- **Both**—In a SPAN session, you can monitor a series or range of ports for both received and sent packets.

Source Port

A source port (also called a *monitored port*) is a switched port that you monitor for network traffic analysis. In a single local SPAN session or RSPAN source session, you can monitor source port traffic such as received (Rx), transmitted (Tx), or bidirectional (both). The switch supports any number of source ports (up to the maximum number of available ports on the switch).

A source port has these characteristics:

- It can be any port type (for example, EtherChannel, Fast Ethernet, Gigabit Ethernet, and so forth).
- It cannot be a destination port.
- Each source port can be configured with a direction (ingress, egress, or both) to monitor. For EtherChannel sources, the monitored direction would apply to all the physical ports in the group.
- Source ports can be in the same or different VLANs.

You can configure a trunk port as a source port. All VLANs active on the trunk are monitored.

Destination Port

Each local SPAN session or RSPAN destination session must have a destination port (also called a *monitoring port*) that receives a copy of traffic from the source port.

The destination port has these characteristics:

- It must reside on the same switch as the source port (for a local SPAN session).
- It can be any Ethernet physical port.
- It cannot be a source port or a reflector port.
- It cannot be an EtherChannel group or a VLAN.

- It can be a physical port that is assigned to an EtherChannel group, even if the EtherChannel group has been specified as a SPAN source. The port is removed from the group while it is configured as a SPAN destination port.
- The port does not transmit any traffic except that required for the SPAN session.
- If ingress traffic forwarding is enabled for a network security device, the destination port forwards traffic at Layer 2.
- It does not participate in spanning tree while the SPAN session is active.
- When it is a destination port, it does not participate in any of the Layer 2 protocols (STP, VTP, CDP, DTP, PagP, or LACP).
- No address learning occurs on the destination port.
- A destination port receives copies of sent and received traffic for all monitored source ports. If a destination port is oversubscribed, it could become congested. This could affect traffic forwarding on one or more of the source ports.

Reflector Port

The reflector port is the mechanism that copies packets onto an RSPAN VLAN. The reflector port forwards only the traffic from the RSPAN source session with which it is affiliated. Any device connected to a port set as a reflector port loses connectivity until the RSPAN source session is disabled.

The reflector port has these characteristics:

- It is a port set to loopback.
- It cannot be an EtherChannel group, it does not trunk, and it cannot do protocol filtering.
- It can be a physical port that is assigned to an EtherChannel group, even if the EtherChannel group is specified as a SPAN source. The port is removed from the group while it is configured as a reflector port.
- A port used as a reflector port cannot be a SPAN source or destination port, nor can a port be a reflector port for more than one session at a time.
- It is invisible to all VLANs.
- The native VLAN for looped-back traffic on a reflector port is the RSPAN VLAN.
- The reflector port loops back untagged traffic to the switch. The traffic is then placed on the RSPAN VLAN and flooded to any trunk ports that carry the RSPAN VLAN.
- Spanning tree is automatically disabled on a reflector port.
- A reflector port receives copies of sent and received traffic for all monitored source ports. If a reflector port is oversubscribed, it could become congested. This could affect traffic forwarding on one or more of the source ports.

If the bandwidth of the reflector port is not sufficient for the traffic volume from the corresponding source ports, the excess packets are dropped. A Gigabit port reflects at 1 Gbps.

SPAN Traffic

You can use local SPAN to monitor all network traffic, including multicast and bridge protocol data unit (BPDU) packets, and Cisco Discovery Protocol (CDP), VLAN Trunk Protocol (VTP), Dynamic Trunking Protocol (DTP), Spanning Tree Protocol (STP), Port Aggregation Protocol (PagP), and Link Aggregation Control Protocol (LACP) packets. You cannot use RSPAN to monitor Layer 2 protocols. See the “RSPAN Configuration Guidelines” section on page 347 for more information.

In some SPAN configurations, multiple copies of the same source packet are sent to the SPAN destination port. For example, a bidirectional (both Rx and Tx) SPAN session is configured for the sources a1 Rx monitor and the a2 Rx and Tx monitor to destination port d1. If a packet enters the switch through a1 and is switched to a2, both incoming and outgoing packets are sent to destination port d1.

SPAN and RSPAN Interaction with Other Features

SPAN interacts with these features:

- Spanning Tree Protocol (STP)—A destination port or a reflector port does not participate in STP while its SPAN or RSPAN session is active. The destination or reflector port can participate in STP after the SPAN or RSPAN session is disabled. On a source port, SPAN does not affect the STP status. STP can be active on trunk ports carrying an RSPAN VLAN.
- Cisco Discovery Protocol (CDP)—A SPAN destination port does not participate in CDP while the SPAN session is active. After the SPAN session is disabled, the port again participates in CDP.
- VLAN Trunking Protocol (VTP)—You can use VTP to prune an RSPAN VLAN between switches.
- VLAN and trunking—You can modify VLAN membership or trunk settings for source, destination, or reflector ports at any time. However, changes in VLAN membership or trunk settings for a destination or reflector port do not take effect until you disable the SPAN or RSPAN session. Changes in VLAN membership or trunk settings for a source port immediately take effect, and the SPAN session automatically adjusts accordingly.
- EtherChannel—You can configure an EtherChannel group as a source port but not as a SPAN destination port. When a group is configured as a SPAN source, the entire group is monitored.

If a port is added to a monitored EtherChannel group, the new port is added to the SPAN source port list. If a port is removed from a monitored EtherChannel group, it is automatically removed from the source port list. If the port is the only port in the EtherChannel group, the EtherChannel group is removed from SPAN.

If a physical port that belongs to an EtherChannel group is configured as a SPAN source, destination, or reflector port, it is removed from the group. After the port is removed from the SPAN session, it rejoins the EtherChannel group. Ports removed from an EtherChannel group remain members of the group, but they are in the *down* or *standalone* state.

If a physical port that belongs to an EtherChannel group is a destination or reflector port and the EtherChannel group is a source, the port is removed from the EtherChannel group and from the list of monitored ports.

- QoS—For ingress monitoring, the packets sent to the SPAN destination port might be different from the packets actually received at the SPAN source port because the packets are forwarded after ingress QoS classification and policing. The packet DSCP might not be the same as the received packet.
- Multicast traffic can be monitored. For egress and ingress port monitoring, only a single unedited packet is sent to the SPAN destination port. It does not reflect the number of times the multicast packet is sent.
- Port security—A secure port cannot be a SPAN destination port.

SPAN and RSPAN Session Limits

You can configure (and store in NVRAM) one local SPAN session or multiple RSPAN sessions on a switch. The number of active sessions and combinations are subject to these restrictions:

- SPAN or RSPAN source (rx, tx, both): 1 active session limit. (SPAN and RSPAN are mutually exclusive on a source switch).
- RSPAN source sessions have one destination per session with an RSPAN VLAN associated for that session.
- Each RSPAN destination session has one or more destination interfaces for each RSPAN VLAN that they support.
- RSPAN destination sessions are limited to two, or one if a local SPAN or a source RSPAN session is configured on the same switch.

Default SPAN and RSPAN Configuration

Table 69 shows the default SPAN and RSPAN configuration.

Table 69. Default SPAN and RSPAN Configuration

Feature	Default Setting
SPAN state	Disabled.
Source port traffic to monitor	Both received and sent traffic (both).
Encapsulation type (destination port)	Native form (no encapsulation type header).
Ingress forwarding (destination port)	Disabled.

Configuring SPAN

This section describes how to configure SPAN on your switch. It contains this configuration information:

- SPAN Configuration Guidelines, on page 342
- Creating a SPAN Session and Specifying Ports to Monitor, on page 343
- Creating a SPAN Session and Enabling Ingress Traffic, on page 344
- Removing Ports from a SPAN Session, on page 346

SPAN Configuration Guidelines

Follow these guidelines when configuring SPAN:

- SPAN sessions can coexist with RSPAN sessions within the limits described in the “SPAN and RSPAN Session Limits” section on page 342.
- The destination port cannot be a source port; a source port cannot be a destination port.
- You can have only one destination port.
- An EtherChannel port can be a SPAN source port; it cannot be a SPAN destination port.
- For SPAN source ports, you can monitor sent and received traffic for a single port or for a series or range of ports.

- When you configure a switch port as a SPAN destination port, it is no longer a normal switch port; only monitored traffic passes through the SPAN destination port.
- You can configure a disabled port to be a source or destination port, but the SPAN function does not start until the destination port and at least one source port is enabled.
- A SPAN destination port never participates in any VLAN spanning tree. SPAN does include BPDUs in the monitored traffic, so any spanning-tree BPDUs received on the SPAN destination port for a SPAN session were copied from the SPAN source ports.
- When SPAN is enabled, configuration changes have these results:
 - If you change the VLAN configuration of a destination port, the change is not effective until SPAN is disabled.
 - If you disable all source ports or the destination port, the SPAN function stops until both a source and the destination port are enabled.

Creating a SPAN Session and Specifying Ports to Monitor

Beginning in privileged EXEC mode, follow these steps to create a SPAN session and specify the source (monitored) and destination (monitoring) ports:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	no monitor session { <i>session_number</i> all local remote }	Clear any existing SPAN configuration for the session. For <i>session_number</i> , specify 1. Specify all to remove all SPAN sessions, local to remove all local sessions, or remote to remove all remote SPAN sessions.
3.	monitor session <i>session_number</i> source interface <i>interface-id</i> [, -] [both rx tx]	Specify the SPAN session and the source port (monitored port). For <i>session_number</i> , specify 1. For <i>interface-id</i> , specify the source port to monitor. Valid interfaces include physical interfaces and port-channel logical interfaces (port-channel <i>port-channel-number</i>). (Optional) [, -] Specify a series or range of interfaces. Enter a space before and after the comma; enter a space before and after the hyphen. (Optional) Specify the direction of traffic to monitor. If you do not specify a traffic direction, the source interface sends both sent and received traffic. <ul style="list-style-type: none"> • both—Monitor both received and sent traffic. • rx—Monitor received traffic. • tx—Monitor sent traffic.

Step	Command	Purpose
4.	monitor session <i>session_number</i> destination interface <i>interface-id</i> [encapsulation {dot1q}]	Specify the SPAN session and the destination port (monitoring port). For <i>session_number</i> , specify 1. For <i>interface-id</i> , specify the destination port. Valid interfaces include physical interfaces. (Optional) Specify the encapsulation header for outgoing packets. If not specified, packets are sent in native form. <ul style="list-style-type: none"> dot1q—Use 802.1Q encapsulation.
5.	end	Return to privileged EXEC mode.
6.	show monitor [session <i>session_number</i>]	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to set up a SPAN session, session 1, for monitoring source port traffic to a destination port. First, any existing SPAN configuration for session 1 is cleared, and then bidirectional traffic is mirrored from source port 8 to destination port 20.

```
Switch(config)# no monitor session 1
Switch(config)# monitor session 1 source interface gigabitethernet0/8
Switch(config)# monitor session 1 destination interface gigabitethernet0/20
encapsulation dot1q
Switch(config)# end
```

Creating a SPAN Session and Enabling Ingress Traffic

Beginning in privileged EXEC mode, follow these steps to create a SPAN session, to specify the source and destination ports, and to enable ingress traffic on the destination port for a network security device (such as a Cisco IDS Sensor Appliance):

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	no monitor session { <i>session_number</i> all local remote }	Clear any existing SPAN configuration for the session. For <i>session_number</i> , specify 1. Specify all to remove all SPAN sessions, local to remove all local sessions, or remote to remove all remote SPAN sessions.

Step	Command	Purpose
3.	monitor session <i>session_number</i> source interface <i>interface-id</i> [, -] [both rx tx]	Specify the SPAN session and the source port (monitored port). For <i>session_number</i> , specify 1. For <i>interface-id</i> , specify the source port to monitor. Valid interfaces include physical interfaces and port-channel logical interfaces (port-channel <i>port-channel-number</i>). (Optional) [, -] Specify a series or range of interfaces. Enter a space before and after the comma; enter a space before and after the hyphen. (Optional) Specify the direction of traffic to monitor. If you do not specify a traffic direction, the source interface sends both sent and received traffic. <ul style="list-style-type: none"> • both—Monitor both received and sent traffic. • rx—Monitor received traffic. • tx—Monitor sent traffic.
4.	monitor session <i>session_number</i> destination interface <i>interface-id</i> [encapsulation {dot1q}] [ingress vlan <i>vlan id</i>]	Specify the SPAN session, the destination port (monitoring port), the packet encapsulation, and the ingress VLAN. For <i>session_number</i> , specify 1. For <i>interface-id</i> , specify the destination port. Valid interfaces include physical interfaces. (Optional) Specify the encapsulation header for outgoing packets. If not specified, packets are sent in native form. <ul style="list-style-type: none"> • dot1q—Use 802.1Q encapsulation. (Optional) Enter ingress vlan <i>vlan id</i> to enable ingress forwarding and specify a default VLAN.
5.	end	Return to privileged EXEC mode.
6.	show monitor [session <i>session_number</i>]	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to configure the destination port for ingress traffic on VLAN 5 by using a security device that does not support 802.1Q encapsulation.

```
Switch(config)# monitor session 1 destination interface Gi0/17 ingress vlan 5
```

This example shows how to configure the destination port for ingress traffic on VLAN 5 by using a security device that supports 802.1Q encapsulation.

```
Switch(config)# monitor session 1 destination interface Gi0/17 encapsulation dot1q  
ingress vlan 5
```

This example shows how to disable ingress traffic forwarding on the destination port.

```
Switch(config)# monitor session 1 destination interface Gi 0/17
encapsulation dot1q
```

Removing Ports from a SPAN Session

Beginning in privileged EXEC mode, follow these steps to remove a port as a SPAN source for a session:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	no monitor session <i>session_number</i> source interface <i>interface-id</i> [, -] [both rx tx]	Specify the characteristics of the source port (monitored port) and SPAN session to remove. For <i>session</i> , specify 1. For <i>interface-id</i> , specify the source port to no longer monitor. Valid interfaces include physical interfaces and port-channel logical interfaces (port-channel <i>port-channel-number</i>). (Optional) Use [, -] to specify a series or range of interfaces if they were configured. This option is valid when monitoring only received traffic. Enter a space before and after the comma; enter a space before and after the hyphen. (Optional) Specify the direction of traffic (both , rx , or tx) to no longer monitor. If you do not specify a traffic direction, both transmit and receive are disabled.
3.	end	Return to privileged EXEC mode.
4.	show monitor [session <i>session_number</i>]	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a source or destination port from the SPAN session, use the **no monitor session** *session_number* **source interface** *interface-id* global configuration command or the **no monitor session** *session_number* **destination interface** *interface-id* global configuration command. To change the encapsulation type back to the default (native), use the **monitor session** *session_number* **destination interface** *interface-id* without the **encapsulation** keyword.

This example shows how to remove port 17 as a SPAN source for SPAN session 1:

```
Switch(config)# no monitor session 1 source interface gigabitethernet0/17
Switch(config)# end
```

This example shows how to disable received traffic monitoring on port 17, which was configured for bidirectional monitoring:

```
Switch(config)# no monitor session 1 source interface gigabitethernet0/17 rx
```

The monitoring of traffic received on port 1 is disabled, but traffic sent from this port continues to be monitored.

Configuring RSPAN

This section describes how to configure RSPAN on your switch. It contains this configuration information:

- RSPAN Configuration Guidelines, on page 347
- Creating an RSPAN Session, on page 348
- Creating an RSPAN Destination Session, on page 349
- Removing Ports from an RSPAN Session, on page 350

RSPAN Configuration Guidelines

Follow these guidelines when configuring RSPAN:

- All the items in the “SPAN Configuration Guidelines” section on page 342 apply to RSPAN.

Note: As RSPAN VLANs have special properties, you should reserve a few VLANs across your network for use as RSPAN VLANs; do not assign access ports to these VLANs.

- RSPAN sessions can coexist with SPAN sessions within the limits described in the “SPAN and RSPAN Session Limits” section on page 342.
- For RSPAN configuration, you can distribute the source ports and the destination ports across multiple switches in your network.
- A port cannot serve as an RSPAN source port or RSPAN destination port while designated as an RSPAN reflector port.
- When you configure a switch port as a reflector port, it is no longer a normal switch port; only looped-back traffic passes through the reflector port.
- RSPAN does not support BPDU packet monitoring or other Layer 2 switch protocols.
- In a network consisting of only switches, you must use a unique RSPAN VLAN session on each source switch. If more than one source switch uses the same RSPAN VLAN, the switches are limited to act only as source switches to ensure the delivery of all monitored traffic to the destination switch.
- You can configure any VLAN as an RSPAN VLAN as long as these conditions are met:
 - The RSPAN VLAN is not configured as a native VLAN.
 - Extended range RSPAN VLANs will not be propagated to other switches using VTP.
 - No access port is configured in the RSPAN VLAN.
 - All participating switches support RSPAN.

Note: The RSPAN VLAN cannot be VLAN 1 (the default VLAN) or VLAN IDs 1002 through 1005 (reserved to Token Ring and FDDI VLANs).

- You should create an RSPAN VLAN before configuring an RSPAN source or destination session.
- If you enable VTP and VTP pruning, RSPAN traffic is pruned in the trunks to prevent the unwanted flooding of RSPAN traffic across the network for VLAN-IDs that are lower than 1005.

Creating an RSPAN Session

First create an RSPAN VLAN that *does not* exist for the RSPAN session in any of the switches that will participate in RSPAN. With VTP enabled in the network, you can create the RSPAN VLAN in one switch, and VTP propagates it to the other switches in the VTP domain for VLAN-IDs that are lower than 1005. See the “Creating or Modifying an Ethernet VLAN” section on page 249 for more information about creating an RSPAN VLAN.

Use VTP pruning to get efficient flow of RSPAN traffic, or manually delete the RSPAN VLAN from all trunks that do not need to carry the RSPAN traffic.

After creating the RSPAN VLAN, begin in privileged EXEC mode, and follow these steps to start an RSPAN source session and to specify the source (monitored) ports and the destination RSPAN VLAN.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	no monitor session { <i>session_number</i> all local remote }	<p>Clear any existing RSPAN configuration for the session.</p> <p>For <i>session_number</i>, specify the session number identified with this RSPAN session.</p> <p>Specify all to remove all RSPAN sessions, local to remove all local sessions, or remote to remove all remote SPAN sessions.</p>
3.	monitor session <i>session_number</i> source interface <i>interface-id</i> [, -] [both rx tx]	<p>Specify the RSPAN session and the source port (monitored port).</p> <p>For <i>session_number</i>, specify the session number identified with this RSPAN session.</p> <p>For <i>interface-id</i>, specify the source port to monitor. Valid interfaces include physical interfaces and port-channel logical interfaces (port-channel <i>port-channel-number</i>).</p> <p>(Optional) [, -] Specify a series or range of interfaces. Enter a space before and after the comma; enter a space before and after the hyphen.</p> <p>(Optional) Specify the direction of traffic to monitor. If you do not specify a traffic direction, the source interface sends both sent and received traffic.</p> <ul style="list-style-type: none"> • both—Monitor both received and sent traffic. • rx—Monitor received traffic. • tx—Monitor sent traffic.

Step	Command	Purpose
4.	monitor session <i>session_number</i> destination remote vlan <i>vlan-id</i> reflector-port <i>interface</i>	Specify the RSPAN session, the destination remote VLAN, and the reflector port. For <i>session_number</i> , enter the session number identified with this RSPAN session. For <i>vlan-id</i> , specify the RSPAN VLAN to carry the monitored traffic to the destination port. (See the “Creating or Modifying an Ethernet VLAN” section on page 249 for more information about creating an RSPAN VLAN.) For <i>interface</i> , specify the interface that will flood the RSPAN traffic onto the RSPAN VLAN.
5.	end	Return to privileged EXEC mode.
6.	show monitor [session <i>session_number</i>]	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to clear any existing RSPAN configuration for session 1, configure RSPAN session 1 to monitor multiple source interfaces, and configure the destination RSPAN VLAN and the reflector-port.

```
Switch(config)# no monitor session 1
Switch(config)# monitor session 1 source interface gigabitethernet0/17 tx
Switch(config)# monitor session 1 source interface gigabitethernet0/18 rx
Switch(config)# monitor session 1 source interface gigabitethernet0/19 rx
Switch(config)# monitor session 1 source interface port-channel 102 rx
Switch(config)# monitor session 1 destination remote vlan 901 reflector-port
gigabitethernet0/20
Switch(config)# end
```

Creating an RSPAN Destination Session

Beginning in privileged EXEC mode, follow these steps to create an RSPAN destination session and to specify the source RSPAN VLAN and the destination port:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	monitor session <i>session_number</i> source remote vlan <i>vlan-id</i>	Specify the RSPAN session and the source RSPAN VLAN. For <i>session_number</i> , specify the session number identified with this RSPAN session. For <i>vlan-id</i> , specify the source RSPAN VLAN to monitor.
3.	monitor session <i>session_number</i> destination interface <i>interface-id</i> [encapsulation {dot1q}]	Specify the RSPAN session and the destination interface. For <i>session_number</i> , specify the session number identified with this RSPAN session. For <i>interface-id</i> , specify the destination interface. (Optional) Specify the encapsulation header for outgoing packets. If not specified, packets are sent in native form. • dot1q —Use 802.1Q encapsulation.
4.	end	Return to privileged EXEC mode.

Step	Command	Purpose
5.	show monitor [session <i>session_number</i>]	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to configure VLAN 901 as the source remote VLAN and port 17 as the destination interface:

```
Switch(config)# monitor session 1 source remote vlan 901
Switch(config)# monitor session 1 destination interface gigabitethernet0/17
Switch(config)# end
```

Removing Ports from an RSPAN Session

Beginning in privileged EXEC mode, follow these steps to remove a port as an RSPAN source for a session:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	no monitor session <i>session_number</i> source interface <i>interface-id</i> [, -] [both rx tx]	Specify the characteristics of the RSPAN source port (monitored port) to remove. For <i>session_number</i> , specify the session number identified with this RSPAN session. For <i>interface-id</i> , specify the source port to no longer monitor. Valid interfaces include physical interfaces and port-channel logical interfaces (port-channel <i>port-channel-number</i>). (Optional) Use [, -] to specify a series or range of interfaces if they were configured. Enter a space before and after the comma; enter a space before and after the hyphen. (Optional) Specify the direction of traffic (both , rx , or tx) to no longer monitor. If you do not specify a traffic direction, both transmit and receive are disabled.
3.	end	Return to privileged EXEC mode.
4.	show monitor [session <i>session_number</i>]	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to remove port 17 as an RSPAN source for RSPAN session 1:

```
Switch(config)# no monitor session 1 source interface gigabitethernet0/17
Switch(config)# end
```

This example shows how to disable received traffic monitoring on port 17, which was configured for bidirectional monitoring:

```
Switch(config)# no monitor session 1 source interface gigabitethernet0/17 rx
```


The monitoring of traffic received on port 1 is disabled, but traffic sent from this port continues to be monitored.

Displaying SPAN and RSPAN Status

To display the status of the current SPAN or RSPAN configuration, use the **show monitor** privileged EXEC command.

This is an example of output for the **show monitor** privileged EXEC command for SPAN source session 1:

```
Switch# show monitor session 1
Session 1
-----
Type                : Local Session
Source Ports        :
  RX Only           : None
  TX Only           : None
  Both              : Gi0/17
Source VLANs        :
  RX Only           : None
  TX Only           : None
  Both              : None
Source RSPAN VLAN   : None
Destination Ports   : Gi0/20
  Encapsulation     : DOT1Q
  Ingress           : Enabled, default VLAN = 5
Reflector Port      : None
Filter VLANs        : None
Dest RSPAN VLAN     : None
```

Chapter 20. Configuring RMON

This chapter describes how to configure Remote Network Monitoring (RMON) on your switch. RMON is a standard monitoring specification that defines a set of statistics and functions that can be exchanged between RMON-compliant console systems and network probes. RMON provides you with comprehensive network-fault diagnosis, planning, and performance-tuning information.

Note: For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference*.


This chapter consists of these sections:

- Understanding RMON, on page 353
- Configuring RMON, on page 354
- Displaying RMON Status, on page 357

Understanding RMON

RMON is an Internet Engineering Task Force (IETF) standard monitoring specification that allows various network agents and console systems to exchange network monitoring data. You can use the RMON feature with the Simple Network Management Protocol (SNMP) agent in the switch to monitor all the traffic flowing among switches on all connected LAN segments.

Figure 62. Remote Monitoring Example



The switch supports these RMON groups (defined in RFC 1757):

- Statistics (RMON group 1)—Collects Ethernet, Fast Ethernet, and Gigabit Ethernet statistics on an interface.
- History (RMON group 2)—Collects a history group of statistics on Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces for a specified polling interval.
- Alarm (RMON group 3)—Monitors a specific management information base (MIB) object for a specified interval, triggers an alarm at a specified value (rising

threshold), and resets the alarm at another value (falling threshold). Alarms can be used with events; the alarm triggers an event, which can generate a log entry or an SNMP trap.

- Event (RMON group 9)—Determines the action to take when an event is triggered by an alarm. The action can be to generate a log entry or an SNMP trap.

Because switches supported by this release use hardware counters for RMON data processing, the monitoring is more efficient, and little processing power is required.

Configuring RMON

This section describes how to configure RMON on your switch. It contains this configuration information:

- Default RMON Configuration, on page 354
- Configuring RMON Alarms and Events, on page 354
- Configuring RMON Collection on an Interface, on page 356

Default RMON Configuration

RMON is disabled by default; no alarms or events are configured.

Only RMON 1 is supported on the switch.

Configuring RMON Alarms and Events

You can configure your switch for RMON by using the command-line interface (CLI) or an SNMP-compatible network management station. We recommend that you use a generic RMON console application on the network management station (NMS) to take advantage of RMON's network management capabilities. You must also configure SNMP on the switch to access RMON MIB objects. For more information, see Chapter 23 "Configuring SNMP:"

Beginning in privileged EXEC mode, follow these steps to enable RMON alarms and events:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	rmon alarm <i>number variable interval</i> { absolute delta } rising-threshold <i>value</i> [<i>event-number</i>] falling-threshold <i>value</i> [<i>event-number</i>] [<i>owner string</i>]	Set an alarm on a MIB object. <ul style="list-style-type: none"> For <i>number</i>, specify the alarm number. The range is 1 to 65535. For <i>variable</i>, specify the MIB object to monitor. For <i>interval</i>, specify the time in seconds the alarm monitors the MIB variable. The range is 1 to 4294967295 seconds. Specify the absolute keyword to test each MIB variable directly; specify the delta keyword to test the change between samples of a MIB variable. For <i>value</i>, specify a number at which the alarm is triggered and one for when the alarm is reset. The range for the rising threshold and falling threshold <i>values</i> is -2147483648 to 2147483647. (Optional) For <i>event-number</i>, specify the event number to trigger when the rising or falling threshold exceeds its limit. (Optional) For owner string, specify the owner of the alarm.
3.	rmon event <i>number</i> [description <i>string</i>] [log] [<i>owner string</i>] [trap <i>community</i>]	Add an event in the RMON event table that is associated with an RMON event number. <ul style="list-style-type: none"> For <i>number</i>, assign an event number. The range is 1 to 65535. (Optional) For description string, specify a description of the event. (Optional) Use the log keyword to generate an RMON log entry when the event is triggered. (Optional) For owner string, specify the owner of this event. (Optional) For <i>community</i>, enter the SNMP community string used for this trap.
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable an alarm, use the **no rmon alarm** *number* global configuration command on each alarm you configured. You cannot disable at once all the alarms that you configured. To disable an event, use the **no rmon event** *number* global configuration command. To learn more about alarms and events and how they interact with each other, refer to RFC 1757.

You can set an alarm on any MIB object. The following example configures RMON alarm number 10 by using the **rmon alarm** command. The alarm monitors the MIB

variable *ifEntry.20.1* once every 20 seconds until the alarm is disabled and checks the change in the variable's rise or fall. If the *ifEntry.20.1* value shows a MIB counter increase of 15 or more, such as from 100000 to 100015, the alarm is triggered. The alarm in turn triggers event number 1, which is configured with the **rmon event** command. Possible events can include a log entry or an SNMP trap. If the *ifEntry.20.1* value changes by 0, the alarm is reset and can be triggered again.

```
Switch(config)# rmon alarm 10 ifEntry.20.1 20 delta rising-threshold 15 1
falling-threshold 0 owner jjohnson
```

The following example creates RMON event number 1 by using the **rmon event** command. The event is defined as *High ifOutErrors* and generates a log entry when the event is triggered by the alarm. The user *jjones* owns the row that is created in the event table by this command. This example also generates an SNMP trap when the event is triggered.

```
Switch(config)# rmon event 1 log trap eventtrap description "High
ifOutErrors" owner jjones
```

Configuring RMON Collection on an Interface

You must first configure RMON alarms and events to display collection information.

Beginning in privileged EXEC mode, follow these steps to collect group history statistics on an interface:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface on which to collect history.
3.	rmon collection history <i>index</i> [buckets <i>bucket-number</i>] [interval <i>seconds</i>] [owner <i>ownername</i>]	Enable history collection for the specified number of buckets and time period. <ul style="list-style-type: none"> For <i>index</i>, identify the RMON group of statistics. The range is 1 to 65535. (Optional) For buckets <i>bucket-number</i>, specify the maximum number of buckets desired for the RMON collection history group of statistics. The range is 1 to 65535. The default is 50 buckets. (Optional) For interval <i>seconds</i>, specify the number of seconds in each polling cycle. (Optional) For owner <i>ownername</i>, enter the name of the owner of the RMON group of statistics.
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	show rmon history	Display the contents of the switch history table.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable history collection, use the **no rmon collection history** *index* interface configuration command.

Beginning in privileged EXEC mode, follow these steps to collect group Ethernet statistics on an interface:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface on which to collect statistics.
3.	rmon collection stats <i>index</i> [owner <i>ownername</i>]	Enable RMON statistic collection on the interface. <ul style="list-style-type: none"> For <i>index</i>, specify the RMON group of statistics. The range is from 1 to 65535. (Optional) For owner <i>ownername</i>, enter the name of the owner of the RMON group of statistics.
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	show rmon statistics	Display the contents of the switch statistics table.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the collection of group Ethernet statistics, use the **no rmon collection stats** *index* interface configuration command.

Displaying RMON Status

To display the RMON status, use one or more of the privileged EXEC commands in Table 70:

Table 70. Commands for Displaying RMON Status

Command	Purpose
show rmon	Displays general RMON statistics.
show rmon alarms	Displays the RMON alarm table.
show rmon events	Displays the RMON event table.
show rmon history	Displays the RMON history table.
show rmon statistics	Displays the RMON statistics table.

For information about the fields in these displays, refer to the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference*.

Chapter 21. Configuring System Message Logging

This chapter describes how to configure system message logging on your switch.

Note: For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference*

This chapter consists of these sections:

- Understanding System Message Logging, on page 359
- Configuring System Message Logging, on page 359
- Displaying the Logging Configuration, on page 369

Understanding System Message Logging

By default, a switch sends the output from system messages and **debug** privileged EXEC commands to a logging process. The logging process controls the distribution of logging messages to various destinations, such as the logging buffer, terminal lines, or a UNIX syslog server, depending on your configuration. The process also sends messages to the service port.

Note: The syslog format is compatible with 4.3 BSD UNIX.

When the logging process is disabled, messages are sent only to the service port. The messages are sent as they are generated, so message and debug output are interspersed with prompts or output from other commands. Messages are displayed on the service port after the process that generated them has finished.

You can set the severity level of the messages to control the type of messages displayed on the service port and each of the destinations. You can timestamp log messages or set the syslog source address to enhance real-time debugging and management. For information on possible messages, refer to the system message guide for this release.

You can access logged system messages by using the switch command-line interface (CLI) or by saving them to a properly configured syslog server. The switch software saves syslog messages in an internal buffer. You can remotely monitor system messages by accessing the switch through Telnet, through the service port, or by viewing the logs on a syslog server.

Configuring System Message Logging

These sections describe how to configure system message logging:

- System Log Message Format, on page 360
- Default System Message Logging Configuration, on page 361
- Disabling and Enabling Message Logging, on page 361
- Setting the Message Display Destination Device, on page 362
- Synchronizing Log Messages, on page 363
- Enabling and Disabling Timestamps on Log Messages, on page 364
- Enabling and Disabling Sequence Numbers in Log Messages, on page 365
- Defining the Message Severity Level, on page 365
- Limiting Syslog Messages Sent to the History Table and to SNMP, on page 366
- Configuring UNIX Syslog Servers, on page 367

System Log Message Format

System log messages can contain up to 80 characters and a percent sign (%), which follows the optional sequence number or timestamp information, if configured. Messages are displayed in this format:

seq no:timestamp: %facility-severity-MNEMONIC:description

The part of the message preceding the percent sign depends on the setting of the **service sequence-numbers**, **service timestamps log datetime**, **service timestamps log datetime [localtime] [msec] [show-timezone]**, or **service timestamps log uptime** global configuration command.

Table 71 describes the elements of syslog messages.

Table 71. System Log Message Elements

Element	Description
seq no:	Stamps log messages with a sequence number only if the service sequence-numbers global configuration command is configured. For more information, see the “Enabling and Disabling Sequence Numbers in Log Messages” section on page 365.
<i>timestamp</i> formats: <i>mm/dd hh:mm:ss</i> or <i>hh:mm:ss</i> (short uptime) or <i>d h</i> (long uptime)	Date and time of the message or event. This information appears only if the service timestamps log [datetime log] global configuration command is configured. For more information, see the “Enabling and Disabling Timestamps on Log Messages” section on page 364.
<i>facility</i>	The facility to which the message refers (for example, SNMP, SYS, and so forth). For a list of supported facilities, see Table 74 on page 369.
<i>severity</i>	Single-digit code from 0 to 7 that is the severity of the message. For a description of the severity levels, see Table 73 on page 366.
<i>MNEMONIC</i>	Text string that uniquely describes the message.
<i>description</i>	Text string containing detailed information about the event being reported.

This example shows a partial switch system message:

```
00:00:46: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/17, changed state to up
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/20, changed state to up
00:00:48: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1, changed
state to down
00:00:48: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet0/17, changed state to down 2
*Mar  1 18:46:11: %SYS-5-CONFIG_I: Configured from console by vty2
(10.34.195.36)
18:47:02: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
*Mar  1 18:48:50.483 UTC: %SYS-5-CONFIG_I: Configured from console by vty2
(10.34.195.36)
```

Default System Message Logging Configuration

Table 72 shows the default system message logging configuration.

Table 72. Default System Message Logging Configuration

Feature	Default Setting
System message logging to the service port	Enabled
Service port severity	Debugging (and numerically lower levels; see Table 73 on page 366)
Logging buffer size	4096 bytes
Logging history size	1 message
Timestamps	Disabled
Synchronous logging	Disabled
Logging server	Disabled
Syslog server IP address	None configured
Server facility	Local7 (see Table 74 on page 369)
Server severity	Informational (and numerically lower levels; see Table 73 on page 366)

Disabling and Enabling Message Logging

Message logging is enabled by default. It must be enabled to send messages to any destination other than the service port. When enabled, log messages are sent to a logging process, which logs messages to designated locations asynchronously to the processes that generated the messages.

Beginning in privileged EXEC mode, follow these steps to disable message logging:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	no logging on	Disable message logging.
3.	end	Return to privileged EXEC mode.
4.	show running-config or show logging	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Disabling the logging process can slow down the switch because a process must wait until the messages are written to the service port before continuing. When the logging process is disabled, messages are displayed on the service port as soon as they are produced, often appearing in the middle of command output.

The **logging synchronous** global configuration command also affects the display of messages to the service port. When this command is enabled, messages appear only after you press Return. For more information, see the “Synchronizing Log Messages” section on page 363.

To re-enable message logging after it has been disabled, use the **logging on** global configuration command.

Setting the Message Display Destination Device

If message logging is enabled, you can send messages to specific locations.

Beginning in privileged EXEC mode, use one or more of the following commands to specify the locations that receive messages:

Note: The service port is not accessible.

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	logging buffered [<i>size</i>]	Log messages to an internal buffer. The default buffer size is 4096. The range is 4096 to 4294967295 bytes. Note: Do not make the buffer size too large because the switch could run out of memory for other tasks. Use the show memory privileged EXEC command to view the free processor memory on the switch; however, this value is the maximum available, and the buffer size should <i>not</i> be set to this amount.
3.	logging <i>host</i>	Log messages to a UNIX syslog server host. For <i>host</i> , specify the name or IP address of the host to be used as the syslog server. To build a list of syslog servers that receive logging messages, enter this command more than once. For complete syslog server configuration steps, see the “Configuring UNIX Syslog Servers” section on page 367.
4.	logging file flash: <i>filename</i> [<i>max-file-size</i>] [<i>min-file-size</i>] [<i>severity-level-number</i> <i>type</i>]	Store log messages in a file in flash memory. <ul style="list-style-type: none">• For <i>filename</i>, enter the log message filename.• (Optional) For <i>max-file-size</i>, specify the maximum logging file size. The range is 4096 to 2147483647. The default is 4069 bytes.• (Optional) For <i>min-file-size</i>, specify the minimum logging file size. The range is 1024 to 2147483647. The default is 2048 bytes.• (Optional) For <i>severity-level-number</i> <i>type</i>, specify either the logging severity level or the logging type. The severity range is 0 to 7. For a list of logging type keywords, see Table 73 on page 366. By default, the log file receives debugging messages and numerically lower levels.
5.	end	Return to privileged EXEC mode.
6.	terminal monitor	Log messages to a non-service port terminal during the current session. Terminal parameter-setting commands are set locally and do not remain in effect after the session has ended. You must perform this step for each session to see the debugging messages.
7.	show running-config	Verify your entries.
8.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The **logging buffered** global configuration command copies logging messages to an internal buffer. The buffer is circular, so newer messages overwrite older messages

after the buffer is full. To display the messages that are logged in the buffer, use the **show logging** privileged EXEC command. The first message displayed is the oldest message in the buffer. To clear the contents of the buffer, use the **clear logging** privileged EXEC command.

Synchronizing Log Messages

You can configure the system to synchronize unsolicited messages and **debug** privileged EXEC command output with solicited device output and prompts for a specific service port line or virtual terminal line. You can identify the types of messages to be output asynchronously based on the level of severity. You can also determine the maximum number of buffers for storing asynchronous messages for the terminal after which messages are dropped.

Note: The service port is not accessible.

When synchronous logging of unsolicited messages and **debug** command output is enabled, unsolicited device output is displayed on the service port or printed after solicited device output is displayed or printed. Unsolicited messages and **debug** command output is displayed on the service port after the prompt for user input is returned. Therefore, unsolicited messages and **debug** command output are not interspersed with solicited device output and prompts. After the unsolicited messages are displayed, the service port again displays the user prompt.

Beginning in privileged EXEC mode, follow these steps to configure synchronous logging:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	line [console vty] line-number <i>[ending-line-number]</i>	<p>Specify the line to be configured for synchronous logging of messages.</p> <ul style="list-style-type: none"> Use the console keyword for configurations that occur through the switch service port. Use the line vty line-number command to specify which vty lines are to have synchronous logging enabled. You use a vty connection for configurations that occur through a Telnet session. The range of line numbers is from 0 to 15. <p>You can change the setting of all 16 vty lines at once by entering:</p> <p>line vty 0 15</p> <p>Or you can change the setting of the single vty line being used for your current connection. For example, to change the setting for vty line 2, enter:</p> <p>line vty 2</p> <p>When you enter this command, the mode changes to line configuration.</p>

Step	Command	Purpose
3.	logging synchronous [<i>level severity-level all</i>] [<i>limit number-of-buffers</i>]	Enable synchronous logging of messages. <ul style="list-style-type: none"> (Optional) For level severity-level, specify the message severity level. Messages with a severity level equal to or higher than this value are printed asynchronously. Low numbers mean greater severity and high numbers mean lesser severity. The default is 2. (Optional) Specifying level all means that all messages are printed asynchronously regardless of the severity level. (Optional) For limit number-of-buffers, specify the number of buffers to be queued for the terminal after which new messages are dropped. The default is 20.
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable synchronization of unsolicited messages and debug output, use the **no logging synchronous** [*level severity-level | all*] [*limit number-of-buffers*] line configuration command.

Enabling and Disabling Timestamps on Log Messages

By default, log messages are not timestamped.

Beginning in privileged EXEC mode, follow these steps to enable timestamping of log messages:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	service timestamps log uptime or service timestamps log datetime [<i>msec</i>] [<i>localtime</i>] [<i>show-timezone</i>]	Enable log timestamps. The first command enables timestamps on log messages, showing the time since the system was rebooted. The second command enables timestamps on log messages. Depending on the options selected, the timestamp can include the date, time in milliseconds relative to the local time zone, and the time zone name.
3.	end	Return to privileged EXEC mode.
4.	show running-config	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable timestamps for both debug and log messages, use the **no service timestamps** global configuration command.

This example shows part of a logging display with the **service timestamps log datetime** global configuration command enabled:

```
*Mar 1 18:46:11: %SYS-5-CONFIG_I: Configured from console by vty2
(10.34.195.36)
```

This example shows part of a logging display with the **service timestamps log uptime** global configuration command enabled:

```
00:00:46: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up
```

Enabling and Disabling Sequence Numbers in Log Messages

Because there is a chance that more than one log message can have the same timestamp, you can display messages with sequence numbers so that you can unambiguously refer to a single message. By default, sequence numbers in log messages are not displayed.

Beginning in privileged EXEC mode, follow these steps to enable sequence numbers in log messages:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	service sequence-numbers	Enable sequence numbers.
3.	end	Return to privileged EXEC mode.
4.	show running-config	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable sequence numbers, use the **no service sequence-numbers** global configuration command.

This example shows part of a logging display with sequence numbers enabled:

```
000019: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
```

Defining the Message Severity Level

You can limit messages displayed to the selected device by specifying the severity level of the message, which are described in Table 73.

Beginning in privileged EXEC mode, follow these steps to define the message severity level:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	logging console <i>level</i>	Limit messages logged to the service port (Not accessible). By default, the service port receives debugging messages and numerically lower levels (see Table 73 on page 366).
3.	logging monitor <i>level</i>	Limit messages logged to the terminal lines. By default, the terminal receives debugging messages and numerically lower levels (see Table 73 on page 366).
4.	logging trap <i>level</i>	Limit messages logged to the syslog servers. By default, syslog servers receive informational messages and numerically lower levels (see Table 73 on page 366). For complete syslog server configuration steps, see the “Configuring UNIX Syslog Servers” section on page 367.

Step	Command	Purpose
5.	end	Return to privileged EXEC mode.
6.	show running-config or show logging	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Note: Specifying a *level* causes messages at that level and numerically lower levels to be displayed at the destination.

To disable logging to the service port, use the **no logging console** global configuration command. To disable logging to a terminal other than the service port, use the **no logging monitor** global configuration command. To disable logging to syslog servers, use the **no logging trap** global configuration command.

Table 73 describes the *level* keywords. It also lists the corresponding UNIX syslog definitions from the most severe level to the least severe level.

Table 73. Message Logging Level Keywords

Level Keyword	Level	Description	Syslog Definition
emergencies	0	System unstable	LOG_EMERG
alerts	1	Immediate action needed	LOG_ALERT
critical	2	Critical conditions	LOG_CRIT
errors	3	Error conditions	LOG_ERR
warnings	4	Warning conditions	LOG_WARNING
notifications	5	Normal but significant condition	LOG_NOTICE
informational	6	Informational messages only	LOG_INFO
debugging	7	Debugging messages	LOG_DEBUG

The software generates four other categories of messages:

- Error messages about software or hardware malfunctions, displayed at levels **warnings** through **emergencies**. These types of messages mean that the functionality of the switch is affected. For information on how to recover from these malfunctions, refer to the system message guide for this release.
- Output from the **debug** commands, displayed at the **debugging** level. Debug commands are typically used only by technical assistance representatives.
- Interface up or down transitions and system restart messages, displayed at the **notifications** level. This message is only for information; switch functionality is not affected.
- Reload requests and low-process stack messages, displayed at the **informational** level. This message is only for information; switch functionality is not affected.

Limiting Syslog Messages Sent to the History Table and to SNMP

If you enabled syslog message traps to be sent to an SNMP network management station by using the **snmp-server enable trap** global configuration command, you

can change the level of messages sent and stored in the switch history table. You also can change the number of messages that are stored in the history table.

Messages are stored in the history table because SNMP traps are not guaranteed to reach their destination. By default, one message of the level **warning** and numerically lower levels (see Table 73 on page 366) are stored in the history table even if syslog traps are not enabled.

Beginning in privileged EXEC mode, follow these steps to change the level and history table size defaults:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	logging history level¹	Change the default level of syslog messages stored in the history file and sent to the SNMP server. See Table 73 on page 366 for a list of <i>level</i> keywords. By default, warnings , errors , critical , alerts , and emergencies messages are sent.
3.	logging history size number	Specify the number of syslog messages that can be stored in the history table. The default is to store one message. The range is 1 to 500 messages.
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

1. Table 73 lists the level keywords and severity level. For SNMP usage, the severity level values increase by 1. For example, emergencies equal 1, not 0, and critical equals 3, not 2.

When the history table is full (it contains the maximum number of message entries specified with the **logging history size** global configuration command), the oldest message entry is deleted from the table to allow the new message entry to be stored.

To return the logging of syslog messages to the default level, use the **no logging history** global configuration command. To return the number of messages in the history table to the default value, use the **no logging history size** global configuration command.

Configuring UNIX Syslog Servers

The next sections describe how to configure the UNIX server syslog daemon and how to define the UNIX system logging facility.

Logging Messages to a UNIX Syslog Daemon

Before you can send system log messages to a UNIX syslog server, you must configure the syslog daemon on a UNIX server. Log in as root, and perform these steps:

Note: Some recent versions of UNIX syslog daemons no longer accept by default syslog packets from the network. If this is the case with your system, use the UNIX **man syslogd** command to determine what options must be added to or removed from the syslog command line to enable logging of remote syslog messages.

1. Add a line such as the following to the file `/etc/syslog.conf`:

```
local7.debug /usr/adm/logs/switch.log
```

The **local7** keyword specifies the logging facility to be used; see Table 74 on page 369 for information on the facilities. The **debug** keyword specifies the syslog level; see Table 73 on page 366 for information on the severity levels. The syslog daemon sends messages at this level or at a more severe level to the file specified in the next field. The file must already exist, and the syslog daemon must have permission to write to it.

2. Create the log file by entering these commands at the UNIX shell prompt:

```
$ touch /var/log/switch.log
$ chmod 666 /var/log/switch.log
```

3. Make sure the syslog daemon reads the new changes:

```
$ kill -HUP `cat /etc/syslog.pid`
```

For more information, see the **man syslog.conf** and **man syslogd** commands on your UNIX system.

Configuring the UNIX System Logging Facility

When sending system log messages to an external device, you can cause the switch to identify its messages as originating from any of the UNIX syslog facilities.

Beginning in privileged EXEC mode, follow these steps to configure UNIX system facility message logging:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	logging host	Log messages to a UNIX syslog server host by entering its IP address. To build a list of syslog servers that receive logging messages, enter this command more than once.
3.	logging trap level	Limit messages logged to the syslog servers. By default, syslog servers receive informational messages and lower. See Table 73 on page 366 for <i>level</i> keywords.
4.	logging facility facility-type	Configure the syslog facility. See Table 74 on page 369 for <i>facility-type</i> keywords. The default is local7 .
5.	end	Return to privileged EXEC mode.
6.	show running-config	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a syslog server, use the **no logging host** global configuration command, and specify the syslog server IP address. To disable logging to syslog servers, enter the **no logging trap** global configuration command.

Table 74 lists the UNIX system facilities supported by the Cisco IOS software. For more information about these facilities, consult the operator's manual for your UNIX operating system.

Table 74. Logging Facility-Type Keywords

Facility Type Keyword	Description
auth	Authorization system
cron	Cron facility
daemon	System daemon
kern	Kernel
local0-7	Locally defined messages
lpr	Line printer system
mail	Mail system
news	USENET news
sys9	System use
sys10	System use
sys11	System use
sys12	System use
sys13	System use
sys14	System use
syslog	System log
user	User process
uucp	UNIX-to-UNIX copy system

Displaying the Logging Configuration

To display the logging configuration and the contents of the log buffer, use the **show logging** privileged EXEC command. For information about the fields in this display, refer to the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference*.

Chapter 22. Configuring Network Security with ACLs

This chapter describes how to configure network security on a switch by using access control lists (ACLs), which are also referred to in commands and tables as *access lists*.

You can create ACLs for physical interfaces or management interfaces. A management interface is defined as a management VLAN or any traffic that is going directly to the CPU, such as SNMP, Telnet, or web traffic.

Note: An ACLs that applied is to a physical interface has a limitation of one mask, and certain keywords are not supported. For more information, see the “Guidelines for Applying ACLs to Physical Interfaces” section on page 375.

Note: For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release and the *Cisco IOS IP and IP Routing Configuration Guide, Cisco IOS Release 12.1* and the *Cisco IOS IP and IP Routing Command Reference, Cisco IOS Release 12.1*.

This chapter consists of these sections:

- Understanding ACLs, on page 371
- Configuring ACLs, on page 375
- Applying ACLs to Terminal Lines or Physical Interfaces, on page 387
- Displaying ACL Information, on page 388
- Examples for Compiling ACLs, on page 390

You can configure ACLs by using the Cluster Management Suite (CMS) or through the command-line interface (CLI). Refer to the CMS online help for step-by-step configuration procedures through CMS. For information about accessing and using CMS, see Chapter 3 “Getting Started with CMS.”

You can also use the security wizard to filter inbound traffic on the switches. Filtering can be based on network addresses, Transmission Control Protocol (TCP) applications, or User Datagram Protocol (UDP) applications. You can choose whether to drop or to forward packets that meet the filtering criteria. To use this wizard, you must know how the network is designed and how interfaces are used on the filtering device. Refer to the security wizard online help for step-by-step configuration procedures about using this wizard.

Understanding ACLs

Packet filtering can limit network traffic and restrict network use by certain users or devices. ACLs can filter traffic as it passes through a switch and permit or deny packets at specified interfaces. An ACL is a sequential collection of permit and deny conditions that apply to packets. When a packet is received on an interface, the switch compares the fields in the packet against any applied ACLs to verify that the packet has the required permissions to be forwarded, based on the criteria specified in the access lists. The switch tests the packet against the conditions in an access list one by one. The first match determines whether the switch accepts or rejects the packet. Because the switch stops testing conditions after the first match, the order of conditions in the list is critical. If no conditions match, the switch rejects the packet. If there are no restrictions, the switch forwards the packet; otherwise, the switch drops the packet.

You configure access lists on a Layer 2 switch to provide basic security for your network. If you do not configure ACLs, all packets passing through the switch could be allowed onto all parts of the network. You can use ACLs to control which hosts can

access different parts of a network or to decide which types of traffic are forwarded or blocked at switch interfaces. For example, you can allow e-mail traffic to be forwarded but not Telnet traffic. ACLs can be configured to block inbound traffic.

An ACL contains an ordered list of access control entries (ACEs). Each ACE specifies *permit* or *deny* and a set of conditions the packet must satisfy in order to match the ACE. The meaning of *permit* or *deny* depends on the context in which the ACL is used.


The switch supports these types of ACLs on physical interfaces in the inbound direction:

- IP ACLs filter IP, TCP, and UDP traffic.
- Ethernet or MAC ACLs filter Layer 2 traffic.
- MAC extended access lists use source and destination MAC addresses and optional protocol type information for matching operations.
- Standard IP access lists use source addresses for matching operations.
- Extended IP access lists use source and destination addresses and optional protocol type information for matching operations.

The switch examines access lists associated with features configured on a given interface. As packets enter the switch on an interface, ACLs associated with all inbound features configured on that interface are examined.

ACLs permit or deny packet forwarding based on how the packet matches the entries in the ACL. For example, you can use ACLs to allow one host to access a part of a network, but to prevent another host from accessing the same part. In Figure 63, ACLs applied at the switch input allow Host A to access the Human Resources network, but prevent Host B from accessing the same network.

Figure 63. Using ACLs to Control Traffic to a Network



65285

Handling Fragmented and Unfragmented Traffic

IP packets can be fragmented as they cross the network. When this happens, only the fragment containing the beginning of the packet contains the Layer 4 information, such as TCP or UDP port numbers, Internet Control Message Protocol (ICMP) type and code, and so on. All other fragments are missing this information.

Some ACEs do not check Layer 4 information and therefore can be applied to all packet fragments. ACEs that do test Layer 4 information cannot be applied in the standard manner to most of the fragments in a fragmented IP packet. When the fragment contains no Layer 4 information and the ACE tests some Layer 4 information, the matching rules are modified:

- Permit ACEs that check the Layer 3 information in the fragment (including protocol type, such as TCP, UDP, and so on) are considered to match the fragment regardless of what the missing Layer 4 information might have been.
- Deny ACEs that check Layer 4 information never match a fragment unless the fragment contains Layer 4 information.

Consider access list 102, configured with these commands, applied to three fragmented packets:

```
Switch (config)# access-list 102 permit tcp any host 10.1.1.1 eq smtp
Switch (config)# access-list 102 deny tcp any host 10.1.1.2 eq telnet
Switch (config)# access-list 102 deny tcp any any
```

Note: In the first and second ACEs in the examples, the *eq* keyword after the destination address means to test for the TCP-destination-port well-known numbers equaling Simple Mail Transfer Protocol (SMTP) and Telnet, respectively.

- Packet A is a TCP packet from host 10.2.2.2, port 65000, going to host 10.1.1.1 on the SMTP port. If this packet is fragmented, the first fragment matches the first ACE (a permit), as if it were a complete packet because all Layer 4 information is present. The remaining fragments also match the first ACE, even though they do not contain the SMTP port information because the first ACE only checks Layer 3 information when applied to fragments. (The information in this example is that the packet is TCP and that the destination is 10.1.1.1.)
- Packet B is from host 10.2.2.2, port 65001, going to host 10.1.1.2 on the Telnet port. If this packet is fragmented, the first fragment matches the second ACE (a deny) because all Layer 3 and Layer 4 information is present. The remaining fragments in the packet do not match the second ACE because they are missing Layer 4 information.
- Because the first fragment was denied, host 10.1.1.2 cannot reassemble a complete packet, so packet B is effectively denied. However, the later fragments that are permitted will consume bandwidth on the network and the resources of host 10.1.1.2 as it tries to reassemble the packet.
- Fragmented packet C is from host 10.2.2.2, port 65001, going to host 10.1.1.3, port ftp. If this packet is fragmented, the first fragment matches the third ACE (a deny). All other fragments also match the third ACE because that ACE does not check any Layer 4 information and because Layer 3 information in all fragments shows that they are being sent to host 10.1.1.3, and the earlier permit ACEs were checking different hosts.

Understanding Access Control Parameters

Before configuring ACLs on the switches, you must have a thorough understanding of the access control parameters (ACPs). ACPs are referred to as *masks* in the switch CLI commands, output, and CMS.

Each ACE has a mask and a rule. The Classification Field or mask is the field of interest on which you want to perform an action. The specific values associated with a given mask are called *rules*.

Packets can be classified on these Layer 2, Layer 3, and Layer 4 fields:

- Layer 2 fields:

- Source MAC address (Specify all 48 bits.)
- Destination MAC address (Specify all 48 bits.)
- Ethertype (16-bit ethertype field)

You can use any combination or all of these fields simultaneously to define a flow.

- Layer 3 fields:
 - IP source address (Specify all 32 IP source address bits to define the flow, or specify an user- defined subnet. There are no restrictions on the IP subnet to be specified.)
 - IP destination address (Specify all 32 IP destination address bits to define the flow, or specify an user-defined subnet. There are no restrictions on the IP subnet to be specified.)

You can use any combination or all of these fields simultaneously to define a flow.

- Layer 4 fields:
 - TCP (You can specify a TCP source, destination port number, or both at the same time.)
 - UDP (You can specify a UDP source, destination port number, or both at the same time.)

Note: A mask can be a combination of either multiple Layer 3 and Layer 4 fields or of multiple Layer 2 fields. Layer 2 fields cannot be combined with Layer 3 or Layer 4 fields.

There are two types of masks:

- User-defined mask—masks that are defined by the user.
- System-defined mask—these masks can be configured on any interface:

```
Switch (config-ext-nacl)# permit tcp any any
Switch (config-ext-nacl)# deny tcp any any
Switch (config-ext-nacl)# permit udp any any
Switch (config-ext-nacl)# deny udp any any
Switch (config-ext-nacl)# permit ip any any
Switch (config-ext-nacl)# deny ip any any
Switch (config-ext-nacl)# deny any any
Switch (config-ext-nacl)# permit any any
```

Note: In an IP extended ACL (both named and numbered), a Layer 4 system-defined mask cannot precede a Layer 3 user-defined mask. For example, a Layer 4 system-defined mask such as **permit tcp any any** or **deny udp any any** cannot precede a Layer 3 user-defined mask such as **permit ip 10.1.1.1 any**. If you configure this combination, the ACL is not allowed on a Layer 2 interface. All other combinations of system-defined and user-defined masks are allowed in security ACLs.

The switch ACL configuration is consistent with other Cisco Catalyst switches. However, there are significant restrictions for configuring ACLs on the switches.

Only four user-defined masks can be defined for the entire system. These can be used for either security or quality of service (QoS) but cannot be shared by QoS and security. You can configure as many ACLs as you require. However, a system error message appears if ACLs with more than four different masks are applied to interfaces. For more information about error messages, see the *system message guide for this release*.

Table 75 lists a summary of the ACL restrictions on the switches.

Table 75. Summary of ACL Restrictions

Restriction	Number Permitted
Number of user-defined masks allowed in an ACL	1
Number of ACLs allowed on an interface	1
Total number of user-defined masks for security and QoS allowed on a switch	4
Number of rules allowed per mask	16

Guidelines for Applying ACLs to Physical Interfaces

When applying ACLs to physical interfaces, follow these configuration guidelines:

- Only one ACL can be attached to an interface. For more information, refer to the **ip access-group** interface command in the *command reference* for this release.
- All ACEs in an ACL must have the same user-defined mask. However, ACEs can have different rules that use the same mask. On a given interface, only one type of user-defined mask is allowed, but you can apply any number of system-defined masks. For more information on system-defined masks, see the “Understanding Access Control Parameters” section on page 373.

This example shows the same mask in an ACL:

```
Switch (config)#ip access-list extended acl2
Switch (config-ext-nacl)# permit tcp 10.1.1.1 0.0.0.0 any eq 80
Switch (config-ext-nacl)# permit tcp 20.1.1.1 0.0.0.0 any eq 23
```

In this example, the first ACE permits all the TCP packets coming from host 10.1.1.1 with a destination TCP port number of 80. The second ACE permits all TCP packets coming from host 20.1.1.1 with a destination TCP port number of 23. Both the ACEs use the same mask; therefore, a switch supports this ACL.

- When you apply an ACL to a physical interface, some keywords are not supported and certain mask restrictions apply to the ACLs. See the “Creating a Numbered Standard ACL” section on page 377 and the “Creating a Numbered Extended ACL” section on page 378 for creating these ACLs.

Note: You can also apply ACLs to a management interface without the above limitations. For information, refer to the *Cisco IOS IP and IP Routing Configuration Guide, Cisco IOS Release 12.1* and the *Cisco IOS IP and IP Routing Command Reference, Cisco IOS Release 12.1*.

Configuring ACLs

This section includes these topics:

- “Unsupported Features” section on page 376
- “Creating Standard and Extended IP ACLs” section on page 376
- “Creating Named MAC Extended ACLs” section on page 386
- “Creating MAC Access Groups” section on page 387

Configuring ACLs on a Layer 2 interface is the same as configuring ACLs on Cisco routers. The process is briefly described here. For more detailed information about configuring router ACLs, refer to the *Cisco IP and IP Routing Configuration Guide, Cisco IOS Release 12.1*. For detailed information about the commands, refer to the *Cisco IOS IP and IP Routing Command Reference, Cisco IOS Release 12.1*. For a list

of Cisco IOS features not supported on the switch, see the “Unsupported Features” section on page 376.

Unsupported Features

The switch does not support these Cisco IOS router ACL-related features:

- Non-IP protocol ACLs (see Table 76 on page 377)
- Bridge-group ACLs
- IP accounting
- ACL support on the outbound direction
- Inbound and outbound rate limiting (except with QoS ACLs)
- IP packets that have a header length of less than 5 bytes
- Reflexive ACLs
- Dynamic ACLs (except for certain specialized dynamic ACLs used by the switch clustering feature)
- ICMP-based filtering
- Interior Gateway Routing Protocol (IGMP)-based filtering

Creating Standard and Extended IP ACLs

This section describes how to create switch IP ACLs. The switch tests packets against the conditions in an access list one by one. The first match determines whether the switch accepts or rejects the packet. Because the switch stops testing conditions after the first match, the order of the conditions is critical. If no conditions match, the switch denies the packet.

Follow these steps to use ACLs:

1. Create an ACL by specifying an access list number or name and access conditions.
2. Apply the ACL to interfaces or terminal lines.

The software supports these kinds of IP access lists:

- Standard IP access lists use source addresses for matching operations.
- Extended IP access lists use source and destination addresses for matching operations and optional protocol-type information for finer granularity of control.

Note: MAC extended access list use source and destination MAC addresses and optional protocol type information for matching operations. For more information, see the “Creating Named MAC Extended ACLs” section on page 386.

The next sections describe access lists and the steps for using them.

ACL Numbers

The number you use to denote your ACL shows the type of access list that you are creating. Table 76 lists the access list number and corresponding type and shows whether or not they are supported by the switch. The switch supports IP standard and IP extended access lists, numbers 1 to 199 and 1300 to 2699.

Table 76. Access List Numbers

ACL Number	Type	Supported
1–99	IP standard access list	Yes
100–199	IP extended access list	Yes
200–299	Protocol type-code access list	No
300–399	DECnet access list	No
400–499	XNS standard access list	No
500–599	XNS extended access list	No
600–699	AppleTalk access list	No
700–799	48-bit MAC address access list	No
800–899	IPX standard access list	No
900–999	IPX extended access list	No
1000–1099	IPX SAP access list	No
1100–1199	Extended 48-bit MAC address access list	No
1200–1299	IPX summary address access list	No
1300–1999	IP standard access list (expanded range)	Yes
2000–2699	IP extended access list (expanded range)	Yes

Note: In addition to numbered standard and extended ACLs, you can also create named standard and extended IP ACLs by using the supported numbers. That is, the name of a standard IP ACL can be 1 to 99; the name of an extended IP ACL can be 100 to 199. The advantage of using named ACLs instead of numbered lists is that you can delete individual entries from a named list.

Creating a Numbered Standard ACL

Note: For information about creating ACLs to apply to a management interface, refer to the *Cisco IOS IP and IP Routing Configuration Guide, Cisco IOS Release 12.1* and the *Cisco IOS IP and IP Routing Command Reference, Cisco IOS Release 12.1*. You can these apply these ACLs only to a management interface.

Beginning in privileged EXEC mode, follow these steps to create a numbered standard IP ACL:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	access-list <i>access-list-number</i> { deny permit remark } { <i>source source-wildcard</i> host <i>source</i> any }	<p>Define a standard IP ACL by using a source address and wildcard.</p> <p>The <i>access-list-number</i> is a decimal number from 1 to 99 or 1300 to 1999.</p> <p>Enter deny or permit to specify whether to deny or permit access if conditions are matched.</p> <p>The <i>source</i> is the source address of the network or host from which the packet is being sent:</p> <ul style="list-style-type: none"> • The 32-bit quantity in dotted-decimal format. • The keyword any as an abbreviation for <i>source</i> and <i>source-wildcard</i> of 0.0.0.0 255.255.255.255. You do not need to enter a source wildcard. • The keyword host as an abbreviation for <i>source</i> and <i>source-wildcard</i> of <i>source</i> 0.0.0.0. <p>(Optional) The <i>source-wildcard</i> applies wildcard bits to the source. (See first bullet item.)</p> <p>Note: The log option is not supported on the switches.</p>
3.	end	Return to privileged EXEC mode.
4.	show access-lists [<i>number</i> <i>name</i>]	Show the access list configuration.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no access-list** *access-list-number* global configuration command to delete the entire ACL. You cannot delete individual ACEs from numbered access lists.

Note: When creating an ACL, remember that, by default, the end of the ACL contains an implicit deny statement for all packets that it did not find a match for before reaching the end. With standard access lists, if you omit the mask from an associated IP host address ACL specification, 0.0.0.0 is assumed to be the mask.

This example shows how to create a standard ACL to deny access to IP host 171.69.198.102, permit access to any others, and display the results.

```
Switch (config)# access-list 2 deny host 171.69.198.102
Switch (config)# access-list 2 permit any
Switch(config)# end
Switch# show access-lists
Standard IP access list 2
    deny 171.69.198.102
    permit any
```

Creating a Numbered Extended ACL

Although standard ACLs use only source addresses for matching, you can use an extended ACL source and destination addresses for matching operations and optional protocol type information for finer granularity of control. Some protocols also have specific parameters and keywords that apply to that protocol.

These IP protocols are supported on physical interfaces (protocol keywords are in parentheses in bold): Internet Protocol (**ip**), Transmission Control Protocol (**tcp**), or User Datagram Protocol (**udp**).

Supported parameters can be grouped into these categories:

- TCP
- UDP

Table 77 lists the possible filtering parameters for ACEs for each protocol type.

Table 77. Filtering Parameter ACEs Supported by Different IP Protocols

Filtering Parameter ¹	TCP	UDP
Layer 3 Parameters:		
IP type of service (ToS) byte ²	–	–
Differentiated Services Code Point (DSCP)	X	X
IP source address	X	X
IP destination address	X	X
Fragments	–	–
TCP or UDP	X	X
Layer 4 Parameters		
Source port operator	X	X
Source port	X	X
Destination port operator	X	X
Destination port	X	X
TCP flag	–	–

1.X in a protocol column means support for the filtering parameter.

2.No support for type of service (ToS) minimize monetary cost bit.

For more details about the specific keywords relative to each protocol, refer to the *Cisco IP and IP Routing Command Reference, Cisco IOS Release 12.1*.

Note: The switch does not support dynamic or reflexive access lists. It also does not support filtering based on the minimize-monetary-cost type of service (ToS) bit.

When creating ACEs in numbered extended access lists, remember that after you create the list, any additions are placed at the end of the list. You cannot reorder the list or selectively add or remove ACEs from a numbered list.

Note: For information about creating ACLs to apply to management interfaces, refer to the *Cisco IOS IP and IP Routing Configuration Guide, Release 12.1* and the *Cisco IOS IP and IP Routing Command Reference, Cisco IOS Release 12.1*. You can apply ACLs only to a management interface or the CPU, such as SNMP, Telnet, or Web traffic.

Beginning in privileged EXEC mode, follow these steps to create an extended ACL:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	access-list <i>access-list-number</i> { deny permit remark } <i>protocol</i> { <i>source source-wildcard</i> host <i>source</i> any } [<i>operator port</i>] { <i>destination destination-wildcard</i> host <i>destination</i> any } [<i>operator port</i>] [dscp <i>dscp-value</i>] [time-range <i>time-range-name</i>]	<p>Define an extended IP access list and the access conditions.</p> <p>The <i>access-list-number</i> is a decimal number from 100 to 199 or 2000 to 2699.</p> <p>Enter deny or permit to specify whether to deny or permit the packet if conditions are matched.</p> <p>For <i>protocol</i>, enter the name or number of an IP protocol: IP, TCP, or UDP. To match any Internet protocol (including TCP and UDP), use the keyword ip.</p> <p>The <i>source</i> is the number of the network or host from which the packet is sent.</p> <p>The <i>source-wildcard</i> applies wildcard bits to the source.</p> <p>The <i>destination</i> is the network or host number to which the packet is sent.</p> <p>Define a destination or source port.</p> <ul style="list-style-type: none"> • The <i>operator</i> can be only eq (equal). • If operator is after <i>source source-wildcard</i>, conditions match when the source port matches the defined port. • If operator is after <i>destination destination-wildcard</i>, conditions match when the destination port matches the defined port. • The <i>port</i> is a decimal number or name of a TCP or UDP port. The number can be from 0 to 65535. • Use TCP port names only for TCP traffic. • Use UDP port names only for UDP traffic. <p>The <i>destination-wildcard</i> applies wildcard bits to the destination.</p> <p><i>Source</i>, <i>source-wildcard</i>, <i>destination</i>, and <i>destination-wildcard</i> can be specified in three ways:</p> <ul style="list-style-type: none"> • The 32-bit quantity in dotted-decimal format. • The keyword any as an abbreviation for <i>source</i> and <i>source-wildcard</i> of 0.0.0.0 255.255.255.255 or any source host. • The keyword host, followed by the 32-bit quantity in dotted-decimal format, as an abbreviation for a single host with <i>source</i> and <i>source-wildcard</i> of <i>source</i> 0.0.0.0. <p>dscp—Enter to match packets with any of the supported 13 DSCP values (0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56), or use the question mark (?) to see a list of available values.</p> <p>The time-range keyword is optional. For an explanation of this keyword, see the “Applying Time Ranges to ACLs” section on page 383.</p>

Step	Command	Purpose
3.	show access-lists [<i>number</i> <i>name</i>]	Verify the access list configuration.
4.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no access-list** *access-list-number* global configuration command to delete the entire access list. You cannot delete individual ACEs from numbered access lists.

This example shows how to create and display an extended access list to deny Telnet access from any host in network 171.69.198.0 to any host in network 172.20.52.0 and permit any others. (The **eq** keyword after the destination address means to test for the TCP destination port number equaling Telnet.)

```
Switch(config)# access-list 102 deny tcp 171.69.198.0 0.0.0.255 172.20.52.0
0.0.0.255 eq telnet
Switch(config)# access-list 102 permit tcp any any
Switch(config)# end
```

```
Switch# show access-lists
Extended IP access list 102
  deny tcp 171.69.198.0 0.0.0.255 172.20.52.0 0.0.0.255 eq telnet
  permit tcp any any
```

After an ACL is created, any additions (possibly entered from the terminal) are placed at the end of the list. You can add ACEs to an ACL, but deleting any ACE deletes the entire ACL.

Note: When creating an ACL, remember that, by default, the end of the access list contains an implicit deny statement for all packets if the access list does not find a match before reaching the end. With standard access lists, if you omit the mask from an associated IP host address ACL specification, 0.0.0.0 is assumed to be the mask.

After creating an ACL, you must apply it to a line or interface, as described in the “Applying ACLs to Terminal Lines or Physical Interfaces” section on page 387.

Creating Named Standard and Extended ACLs

You can identify IP ACLs with an alphanumeric string (a name) rather than a number. You can use named ACLs to configure more IP access lists on a switch than if you use numbered access lists. If you identify your access list with a name rather than a number, the mode and command syntax are slightly different. However, not all commands that use IP access lists accept a named ACL.

Note: The name you give to a standard ACL or extended ACL can also be a number in the supported range of access list numbers. That is, the name of a standard IP ACL can be 1 to 99; the name of an extended IP ACL can be 100 to 199. The advantage of using named ACLs instead of numbered lists is that you can delete individual entries from a named list.

Consider these guidelines and limitations before configuring named ACLs:

- A standard ACL and an extended ACL cannot have the same name.
- Numbered ACLs are also available, as described in the “Creating Standard and Extended IP ACLs” section on page 376.

Beginning in privileged EXEC mode, follow these steps to create a standard named access list using names:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	ip access-list standard {name / access-list-number}	Define a standard IP access list by using a name, and enter access-list configuration mode. Note: The name can be a number from 1 to 99.
3.	deny {source source-wildcard host source any} or permit {source source-wildcard host source any}	In access-list configuration mode, specify one or more conditions denied or permitted to determine if the packet is forwarded or dropped. <ul style="list-style-type: none">host source represents a <i>source</i> and <i>source-wildcard</i> of <i>source</i> 0.0.0.0.any represents a <i>source</i> and <i>source-wildcard</i> of 0.0.0.0 255.255.255.255. Note: The log option is not supported on the switches.
4.	end	Return to privileged EXEC mode.
5.	show access-lists [number name]	Show the access list configuration.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Beginning in privileged EXEC mode, follow these steps to create an extended named ACL using names:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	ip access-list extended {name / access-list-number}	Define an extended IP access list by using a name, and enter access-list configuration mode. Note: The name can be a number from 100 to 199.
3.	{deny permit} protocol {source source-wildcard host source any} [operator port] {destination destination-wildcard host destination any} [operator port] [dscp dscp-value] [time-range time-range-name]	In access-list configuration mode, specify the conditions allowed or denied. <i>See the “Creating a Numbered Extended ACL” section on page 378 for definitions of protocols and other keywords.</i> <ul style="list-style-type: none">host source represents a <i>source</i> and <i>source-wildcard</i> of <i>source</i> 0.0.0.0, and host destination represents a <i>destination</i> and <i>destination-wildcard</i> of <i>destination</i> 0.0.0.0.any represents a <i>source</i> and <i>source-wildcard</i> or <i>destination</i> and <i>destination-wildcard</i> of 0.0.0.0 255.255.255.255. dscp —Enter to match packets with any of the supported 13 DSCP values (0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56), or use the question mark (?) to see a list of available values. The time-range keyword is optional. For an explanation of this keyword, see the “Applying Time Ranges to ACLs” section on page 383.
4.	end	Return to privileged EXEC mode.
5.	show access-lists [number name]	Show the access list configuration.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

When making the standard and extended ACL, remember that, by default, the end of the ACL contains an implicit deny statement for everything if it did not find a match before reaching the end. For standard ACLs, if you omit the mask from an associated IP host address access list specification, 0.0.0.0 is assumed to be the mask.

After you create an ACL, any additions are placed at the end of the list. You cannot selectively add ACEs to a specific ACL. However, you can use **no permit** and **no deny** commands to remove ACEs from a named ACL. This example shows how you can delete individual ACEs from a named ACL:

```
Switch(config)# ip access-list extended border-list
Switch(config-ext-nacl)# no permit ip host 10.1.1.3 any
```

Being able to selectively remove lines from a named ACL is one reason you might use named ACLs instead of numbered ACLs.

After creating an ACL, you must apply it to a line or interface, as described in the “Applying ACLs to Terminal Lines or Physical Interfaces” section on page 387.

Applying Time Ranges to ACLs

You can implement extended ACLs based on the time of day and week by using the **time-range** global configuration command. First, define the name and times of the day and week of the time range, and then reference the time range by name in an ACL to apply restrictions to the access list. You can use the time range to define when the permit or deny statements in the ACL are in effect. The **time-range** keyword and argument are referenced in the named and numbered extended ACL task tables in the “Creating Standard and Extended IP ACLs” section on page 376, and the “Creating Named Standard and Extended ACLs” section on page 381.

These are some of the many benefits of using time ranges:

- You have more control over permitting or denying a user access to resources, such as an application (identified by an IP address mask pair and a port number).
- You can control logging messages. ACL entries can log traffic at certain times of the day, but not constantly. Therefore, you can simply deny access without having to analyze many logs generated during peak hours.

Note: The time range relies on the switch system clock. Therefore, you need a reliable clock source. We recommend that you use Network Time Protocol (NTP) to synchronize the switch clock. For more information, see the “Managing the System Time and Date” section on page 93.

Beginning in privileged EXEC mode, follow these steps to configure a time-range parameter for an ACL:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	time-range <i>time-range-name</i>	Identify the time-range by a meaningful name (for example, <i>workhours</i>), and enter time-range configuration mode. The name cannot contain a space or quotation mark and must begin with a letter.

Step	Command	Purpose
3.	absolute [start <i>time date</i>] [<i>end time date</i>] or periodic <i>day-of-the-week hh:mm</i> to [<i>day-of-the-week</i>] <i>hh:mm</i> or periodic { weekdays weekend daily } <i>hh:mm to hh:mm</i>	Specify when the function it will be applied to is operational. Use some combination of these commands; multiple periodic statements are allowed; only one absolute statement is allowed. If more than one absolute statement is configured, only the one configured last is executed.
4.	end	Return to privileged EXEC mode.
5.	show time-range	Verify the time-range configuration.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a configured time-range, use the **no time-range** *time-range-name* global configuration command.

Repeat the steps if you have multiple items that you want operational at different times.

This example shows how to configure time ranges for *workhours* and for company holidays and how to verify your configuration.

```
Switch(config)# time-range workhours
Switch(config-time-range)# periodic weekdays 8:00 to 12:00
Switch(config-time-range)# periodic weekdays 13:00 to 17:00
Switch(config-time-range)# exit
Switch(config)# time-range new_year_day_2000
Switch(config-time-range)# absolute start 00:00 1 Jan 2000 end 23:59 1 Jan
2000
Switch(config-time-range)# exit
Switch(config)# time-range thanksgiving_2000
Switch(config-time-range)# absolute start 00:00 22 Nov 2000 end 23:59 23 Nov
2000
Switch(config-time-range)# exit
Switch(config)# time-range christmas_2000
Switch(config-time-range)# absolute start 00:00 24 Dec 2000 end 23:50 25 Dec
2000
Switch(config-time-range)# end
Switch# show time-range
time-range entry: christmas_2000 (inactive)
    absolute start 00:00 24 December 2000 end 23:50 25 December 2000
time-range entry: new_year_day_2000 (inactive)
    absolute start 00:00 01 January 2000 end 23:59 01 January 2000
time-range entry: thanksgiving_2000 (inactive)
    absolute start 00:00 22 November 2000 end 23:59 23 November 2000
time-range entry: workhours (inactive)
    periodic weekdays 8:00 to 12:00
    periodic weekdays 13:00 to 17:00
```

To apply a time range, you must reference it by name (for example, *workhours*) in an extended ACL that can implement time ranges. This example shows how to create and verify extended access list 188 that denies TCP traffic from any source to any

destination during the defined holiday time ranges and permits all TCP traffic during work hours.

```
Switch(config)# access-list 188 deny tcp any any time-range
new_year_day_2003
Switch(config)# access-list 188 deny tcp any any time-range
thanksgiving_2003
Switch(config)# access-list 188 deny tcp any any time-range christmas_2003
Switch(config)# access-list 188 permit tcp any any time-range workhours
Switch(config)# end
Switch# show access-lists
Extended IP access list 188
    deny tcp any any time-range new_year_day_2003 (inactive)
    deny tcp any any time-range thanksgiving_2003 (active)
    deny tcp any any time-range christmas_2003 (inactive)
    permit tcp any any time-range workhours (inactive)
```

This example uses named ACLs to permit and deny the same traffic.

```
Switch(config)# ip access-list extended deny_access
Switch(config-ext-nacl)# deny tcp any any time-range new_year_day_2003
Switch(config-ext-nacl)# deny tcp any any time-range thanksgiving_2003
Switch(config-ext-nacl)# deny tcp any any time-range christmas_2003
Switch(config-ext-nacl)# exit
Switch(config)# ip access-list extended may_access
Switch(config-ext-nacl)# permit tcp any any time-range workhours
Switch(config-ext-nacl)# end
Switch# show ip access-lists
Extended IP access list deny_access
    deny tcp any any time-range new_year_day_2003 (inactive)
    deny tcp any any time-range thanksgiving_2003 (inactive)
    deny tcp any any time-range christmas_2003 (inactive)
Extended IP access list may_access
    permit tcp any any time-range workhours (inactive)
```

Including Comments About Entries in ACLs

You can use the **remark** command to include comments (remarks) about entries in any IP standard or extended ACL. The remarks make the ACL easier for you to understand and scan. Each remark line is limited to 100 characters.

The remark can go before or after a permit or deny statement. You should be consistent about where you put the remark so that it is clear which remark describes which permit or deny statement. For example, it would be confusing to have some remarks before the associated permit or deny statements and some remarks after the associated statements.

For IP numbered standard or extended ACLs, use the **access-list** *access-list number* **remark** *remark* global configuration command to include a comment about an access list. To remove the remark, use the **no** form of this command.

In this example, the workstation belonging to Jones is allowed access, and the workstation belonging to Smith is not allowed access:

```
Switch(config)# access-list 1 remark Permit only Jones workstation through
Switch(config)# access-list 1 permit 171.69.2.88
Switch(config)# access-list 1 remark Do not allow Smith workstation through
Switch(config)# access-list 1 deny 171.69.3.13
```

For an entry in a named IP ACL, use the **remark** access-list global configuration command. To remove the remark, use the **no** form of this command.

In this example, the Jones subnet is not allowed to use outbound Telnet:

```
Switch(config)# ip access-list extended telnetting
Switch(config-ext-nacl)# remark Do not allow Jones subnet to telnet out
Switch(config-ext-nacl)# deny tcp host 171.69.2.88 any eq telnet
```

Creating Named MAC Extended ACLs

You can filter Layer 2 traffic on a physical Layer 2 interface by using MAC addresses and named MAC extended ACLs. The procedure is similar to that of configuring other extended named access lists.

Note: Named MAC extended ACLs are used as a part of the **mac access-group** privileged EXEC command.

For more information about the supported non-IP protocols in the **mac access-list extended** command, refer to the command reference for this release.

Note: Matching on any SNAP-encapsulated packet with a nonzero Organizational Unique Identifier (OUI) is not supported.

Beginning in privileged EXEC mode, follow these steps to create a named MAC extended ACL:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	mac access-list extended <i>name</i>	Define an extended MAC access list by using a name.
3.	{deny permit} {any host source MAC address} {any host destination MAC address} [aarp amber appletalk dec-spanning decnet-iv diagnostic dsm etype-6000 etype-8042 lat lavc-sca mop-console mop-dump msdos mumps netbios vines-echo vines-ip xns-idp]	In extended MAC access-list configuration mode, specify to permit or deny any source MAC address or a specific host source MAC address and any destination MAC address. (Optional) You can also enter these options: aarp amber appletalk dec-spanning decnet-iv diagnostic dsm etype-6000 etype-8042 lat lavc-sca mop-console mop-dump msdos mumps netbios vines-echo vines-ip xns-idp —(a non-IP protocol).
4.	end	Return to privileged EXEC mode.
5.	show access-lists [<i>number</i> <i>name</i>]	Show the access list configuration.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no mac access-list extended** *name* global configuration command to delete the entire ACL. You can also delete individual ACEs from named MAC extended ACLs.

This example shows how to create and display an access list named *mac1*, denying only EtherType DECnet Phase IV traffic, but permitting all other types of traffic.

```
Switch(config)# mac access-list extended mac1
Switch(config-ext-macl)# deny any any decnet-iv
Switch(config-ext-macl)# permit any any
Switch(config-ext-macl)# end
Switch # show access-list
Extended MAC access list mac1
```

```
deny any any deernet-iv
permit any any
```

Creating MAC Access Groups

Beginning in privileged EXEC mode, follow these steps to create MAC access groups and to apply a MAC access list to an interface:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Identify a specific interface for configuration, and enter interface configuration mode. The interface must be a Layer 2 interface.
3.	mac access-group { <i>name</i> } { <i>in</i> }	Control access to the specified interface by using the MAC access list name.
4.	end	Return to privileged EXEC mode.
5.	show mac-access group	Display the MAC ACLs applied on the switch.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to apply ACL 2 on Gigabit Ethernet interface 0/17 to filter packets entering the interface:

```
Switch(config)# interface gigabitethernet0/17
Router(config-if)# mac access-group 2 in
```

Note: The **mac access-group** interface configuration command is only valid when applied to a Layer 2 interface.

For inbound ACLs, after receiving a packet, the switch checks the packet against the ACL. If the ACL permits the packet, the switch continues to process the packet. If the ACL rejects the packet, the switch discards the packet. The MAC ACL applies to both IP and non-IP packets.

When you apply an undefined ACL to an interface, the switch acts as if the ACL has not been applied to the interface and permits all packets. Remember this behavior if you use undefined ACLs as a means of network security.

Applying ACLs to Terminal Lines or Physical Interfaces

Note: Before applying an ACL to a physical interface, see the “Guidelines for Applying ACLs to Physical Interfaces” section on page 375.

You can apply ACLs to any management interface. For information on creating ACLs on management interfaces, refer to the *Cisco IOS IP and IP Routing Configuration Guide, Cisco IOS Release 12.1* and the *Cisco IOS IP and IP Routing Command Reference, Cisco IOS Release 12.1*.

Note: The limitations that apply to ACLs on physical interfaces do not apply to ACLs on management interfaces.

After you create an ACL, you can apply it to one or more management interfaces or terminal lines. ACLs can be applied on inbound interfaces. This section describes how to accomplish this task for both terminal lines and network interfaces. Note these guidelines:

- When controlling access to a line, you must use numbered IP ACLs or MAC extended ACLs.
- When controlling access to an interface, you can use named or numbered ACLs.
- Set identical restrictions on all the virtual terminal lines because a user can attempt to connect to any of them.
- If you apply ACLs to a management interface, the ACL only filters packets that are intended for the CPU, such as SNMP, Telnet, or web traffic.
- If you apply ACLs to a management VLAN, see the “Management VLAN” section on page 84.

Applying ACLs to a Physical Interface

Beginning in privileged EXEC mode, follow these steps to control access to a Layer 2 interface:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Identify a specific interface for configuration and enter interface configuration mode. The interface must be a Layer 2 or management interface or a management interface VLAN ID.
3.	ip access-group { <i>access-list-number</i> / <i>name</i> } { <i>in</i> }	Control access to the specified interface.
4.	end	Return to privileged EXEC mode.
5.	show running-config	Display the access list configuration.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to apply access list 2 on Gigabit Ethernet interface 0/20 to filter packets entering the interface:

```
Switch(config)# interface gigabitethernet0/20
Router(config-if)# ip access-group 2 in
```

Note: The **ip access-group** interface configuration command is only valid when applied to a management interface or a Layer 2 physical interface. ACLs cannot be applied to interface port-channels.

For inbound ACLs, after receiving a packet, the switch checks the packet against the ACL. If the ACL permits the packet, the switch continues to process the packet. If the ACL rejects the packet, the switch discards the packet.

When you apply an undefined ACL to an interface, the switch acts as if the ACL has not been applied to the interface and permits all packets. Remember this behavior if you use undefined ACLs for network security.

Displaying ACL Information

You can display the ACLs that are configured on the switch, and you can display the ACLs that have been applied to physical and management interfaces. This section consists of these topics:

- Displaying ACLs, on page 389

- Displaying Access Groups, on page 389

Displaying ACLs

You can display existing ACLs by using **show** commands.

Beginning in privileged EXEC mode, follow these steps to display access lists:

Step	Command	Purpose
1.	show access-lists [<i>number</i> / <i>name</i>]	Show information about all IP and MAC address access lists or about a specific access list (numbered or named).
2.	show ip access-list [<i>number</i> / <i>name</i>]	Show information about all IP address access lists or about a specific IP ACL (numbered or named).

This example shows all standard and extended ACLs:

```
Switch# show access-lists
Standard IP access list 1
    permit 172.20.10.10
Standard IP ACL 10
    permit 12.12.12.12
Standard IP access list 12
    deny 1.3.3.2
Standard IP access list 32
    permit 172.20.20.20
Standard IP access list 34
    permit 10.24.35.56
    permit 23.45.56.34
Extended IP access list 120
Extended MAC access list mac1
```

This example shows only IP standard and extended ACLs.

```
Switch# show ip access-lists
Standard IP access list 1
    permit 172.20.10.10
Standard IP access list 10
    permit 12.12.12.12
Standard IP access list 12
    deny 1.3.3.2
Standard IP access list 32
    permit 172.20.20.20
Standard IP access list 34
    permit 10.24.35.56
    permit 23.45.56.34
Extended IP access list 120
```

Displaying Access Groups

You use the **ip access-group** interface configuration command to apply ACLs to a Layer 3 interface. When IP is enabled on an interface, you can use the **show ip interface *interface-id*** privileged EXEC command to view the input and output access lists on the interface, as well as other interface characteristics. If IP is not enabled on the interface, the access lists are not shown.

This example shows how to view all access groups configured for VLAN 1 and for Gigabit Ethernet interface 0/20:

```
Switch# show ip interface vlan 1
GigabitEthernet0/20 is up, line protocol is down
  Internet address is 10.20.30.1/16
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Outgoing access list is permit Any
  Inbound access list is 13
```

<information truncated>

```
Switch# show ip interface gigabitethernet0/20
GigabitEthernet0/20 is down, line protocol is down
  Inbound access list is ip1
```

The only way to ensure that you can view all configured access groups under all circumstances is to use the **show running-config** privileged EXEC command. To display the ACL configuration of a single interface, use the **show running-config interface *interface-id*** command.

This example shows how to display the ACL configuration of Gigabit Ethernet interface 0/17:

```
Switch# show running-config interface gigabitethernet0/17
Building configuration...

Current configuration :112 bytes
!
interface GigabitEthernet0/17
 ip access-group 11 in
 snmp trap link-status
 no cdp enable
end!
```

Examples for Compiling ACLs


For detailed information about compiling ACLs, refer to the *Security Configuration Guide* and the *Cisco IOS IP and IP Routing Configuration Guide, Cisco IOS Release 12.1*.

Figure 64 shows a small networked office with a stack of switches that are connected to a Cisco router. A host is connected to the network through the Internet using a WAN link.

Use switch ACLs to do these:

- Create a standard ACL, and filter traffic from a specific Internet host with an address 172.20.128.64.
- Create an extended ACL, and filter traffic to deny HTTP access to all Internet hosts but allow all other types of access.

Figure 64. Using Switch ACLs to Control Traffic



This example uses a standard ACL to allow access to a specific Internet host with the address 172.20.128.64.

```
Switch(config)# access-list 6 permit 172.20.128.64 0.0.0.0
Switch(config)# end
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# ip access-group 6 in
```

This example uses an extended ACL to deny traffic from port 80 (HTTP). It permits all other types of traffic.

```
Switch(config)# access-list 106 deny tcp any any eq 80
Switch(config)# access-list 106 permit ip any any
Switch(config)# interface gigabitethernet0/20
Switch(config-if)# ip access-group 106 in
```

Numbered ACL Examples

This example shows that the switch accepts addresses on network 36.0.0.0 subnets and denies all packets coming from 56.0.0.0 subnets. The ACL is then applied to packets entering Gigabit Ethernet interface 0/17.

```
Switch(config)# access-list 2 permit 36.0.0.0 0.255.255.255
Switch(config)# access-list 2 deny 56.0.0.0 0.255.255.255
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# ip access-group 2 in
```

Extended ACL Examples

In this example of using an extended ACL, you have a network connected to the Internet, and you want any host on the network to be able to form TCP Telnet and SMTP connections to any host on the Internet.

```
Switch(config)# access-list 102 permit tcp any 128.88.0.0 0.0.255.255 eq 23
Switch(config)# access-list 102 permit tcp any 128.88.0.0 0.0.255.255 eq 25
Switch(config)# interface gigabitethernet0/17
```

```
Switch(config-if)# ip access-group 102 in
```

SMTP uses TCP port 25 on one end of the connection and a random port number on the other end. The same port numbers are used throughout the life of the connection. Mail packets coming in from the Internet have a destination port of 25. Because the secure system behind the switch always accepts mail connections on port 25, the incoming services are controlled.

Named ACL Example

The Marketing_group ACL allows any TCP Telnet traffic to the destination address and wildcard 171.69.0.0 0.0.255.255 and denies any other TCP traffic. It permits any other IP traffic.

```
Switch(config)# ip access-list extended marketing_group
Switch(config-ext-nacl)# permit tcp any 171.69.0.0 0.0.255.255 eq telnet
Switch(config-ext-nacl)# deny tcp any any
Switch(config-ext-nacl)# permit ip any any
```

The ACLs are applied to permit Gigabit Ethernet port 0/17, which is configured as a Layer 2 port, with the Marketing_group ACL applied to incoming traffic.

```
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# ip access-group marketing_group in
...
```

Commented IP ACL Entry Examples

In this example of a numbered ACL, the workstation belonging to Jones is allowed access, and the workstation belonging to Smith is not allowed access:

```
Switch(config)# access-list 1 remark Permit only Jones workstation through
Switch(config)# access-list 1 permit 171.69.2.88
Switch(config)# access-list 1 remark Do not allow Smith workstation through
Switch(config)# access-list 1 deny 171.69.3.13
```

In this example of a numbered ACL, the Winter and Smith workstations are not allowed to browse the web:

```
Switch(config)# access-list 100 remark Do not allow Winter to browse the web
Switch(config)# access-list 100 deny host 171.69.3.85 any eq www
Switch(config)# access-list 100 remark Do not allow Smith to browse the web
Switch(config)# access-list 100 deny host 171.69.3.13 any eq www
```

In this example of a named ACL, the Jones subnet is not allowed access:

```
Switch(config)# ip access-list standard prevention
Switch(config-std-nacl)# remark Do not allow Jones subnet through
Switch(config-std-nacl)# deny 171.69.0.0 0.0.255.255
```

In this example of a named ACL, the Jones subnet is not allowed to use outbound Telnet:

```
Switch(config)# ip access-list extended telnetting
Switch(config-ext-nacl)# remark Do not allow Jones subnet to telnet out
```

```
Switch(config-ext-nacl)# deny tcp 171.69.0.0 0.0.255.255 any eq telnet
```

Chapter 23. Configuring SNMP

This chapter describes how to configure the Simple Network Management Protocol (SNMP) on your switch.

Note: For complete syntax and usage information for the commands used in this chapter, refer to the switch command reference for this release and to the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference*.

This chapter consists of these sections:

- Understanding SNMP, on page 395
- Configuring SNMP, on page 399
- Displaying SNMP Status, on page 409

Understanding SNMP

SNMP is an application-layer protocol that provides a message format for communication between managers and agents. The SNMP system consists of an SNMP manager, an SNMP agent, and a management information base (MIB). The SNMP manager can be part of a network management system (NMS) such as CiscoWorks. The agent and MIB reside on the switch. To configure SNMP on the switch, you define the relationship between the manager and the agent.

The SNMP agent contains MIB variables whose values the SNMP manager can request or change. A manager can get a value from an agent or store a value into the agent. The agent gathers data from the MIB, the repository for information about device parameters and network data. The agent can also respond to a manager's requests to get or set data.

An agent can send unsolicited traps to the manager. Traps are messages alerting the SNMP manager to a condition on the network. Traps can mean improper user authentication, restarts, link status (up or down), MAC address tracking, closing of a TCP connection, loss of connection to a neighbor, or other significant events.

This section includes information about these topics:

- SNMP Versions, on page 395
- SNMP Manager Functions, on page 397
- SNMP Agent Functions, on page 397
- SNMP Community Strings, on page 397
- Using SNMP to Access MIB Variables, on page 398
- SNMP Notifications, on page 398

SNMP Versions

This software release supports these SNMP versions:

- SNMPv1—The Simple Network Management Protocol, a Full Internet Standard, defined in RFC 1157.
- SNMPv2C replaces the Party-based Administrative and Security Framework of SNMPv2Classic with the community-string-based Administrative Framework of SNMPv2C while retaining the bulk retrieval and improved error handling of SNMPv2Classic. It has these features:

- SNMPv2—Version 2 of the Simple Network Management Protocol, a Draft Internet Standard, defined in RFCs 1902 through 1907.
- SNMPv2C—The community-string-based Administrative Framework for SNMPv2, an Experimental Internet Protocol defined in RFC 1901.
- SNMPv3—Version 3 of the SNMP is an interoperable standards-based protocol defined in RFCs 2273 to 2275. SNMPv3 provides secure access to devices by authenticating and encrypting packets over the network and includes these security features:
 - Message integrity—ensures that a packet was not tampered with in transit
 - Authentication—determines that the message is from a valid source
 - Encryption—mixes the contents of a package to prevent it from being read by an unauthorized source.

Note: To select encryption, enter the **priv** keyword. This keyword is available only when the cryptographic (encrypted) software image is installed.

Both SNMPv1 and SNMPv2C use a community-based form of security. The community of managers able to access the agent's MIB is defined by an IP address access control list and password.

SNMPv2C includes a bulk retrieval mechanism and more detailed error message reporting to management stations. The bulk retrieval mechanism retrieves tables and large quantities of information, minimizing the number of round-trips required. The SNMPv2C improved error-handling includes expanded error codes that distinguish different kinds of error conditions; these conditions are reported through a single error code in SNMPv1. Error return codes in SNMPv2C report the error type.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy set up for a user and the group within which the user resides. A security level is the permitted level of security within a security model. A combination of the security level and the security model determine which security mechanism is used when handling an SNMP packet. Available security models are SNMPv1, SNMPv2C, and SNMPv3.

Table 78 identifies the characteristics of the different combinations of security models and levels.

Table 78. SNMP Security Models and Levels

Model	Level	Authentication	Encryption	Result
SNMPv1	noAuthNoPriv	Community string	No	Uses a community string match for authentication.
SNMPv2C	noAuthNoPriv	Community string	No	Uses a community string match for authentication.
SNMPv3	noAuthNoPriv	Username	No	Uses a username match for authentication.
SNMPv3	authNoPriv	MD5 or SHA	No	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms.
SNMPv3	authPriv (requires the cryptographic software image)	MD5 or SHA	DES	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms. Provides DES 56-bit encryption in addition to authentication based on the CBC-DES (DES-56) standard.

You must configure the SNMP agent to use the SNMP version supported by the management station. Because an agent can communicate with multiple managers,

you can configure the software to support communications with one management station using the SNMPv1 protocol, one using the SNMPv2C protocol and another using SNMPv3.

SNMP Manager Functions

The SNMP manager uses information in the MIB to perform the operations described in Table 79.

Table 79. SNMP Operations

Operation	Description
get-request	Retrieves a value from a specific variable.
get-next-request	Retrieves a value from a variable within a table. ¹
get-bulk-request ²	Retrieves large blocks of data, such as multiple rows in a table, that would otherwise require the transmission of many small blocks of data.
get-response	Replies to a get-request, get-next-request, and set-request sent by an NMS.
set-request	Stores a value in a specific variable.
trap	An unsolicited message sent by an SNMP agent to an SNMP manager when some event has occurred.

1. With this operation, an SNMP manager does not need to know the exact variable name. A sequential search is performed to find the needed variable from within a table.

2. The **get-bulk** command only works with SNMPv2 or later.

SNMP Agent Functions

The SNMP agent responds to SNMP manager requests as follows:

- Get a MIB variable—The SNMP agent begins this function in response to a request from the NMS. The agent retrieves the value of the requested MIB variable and responds to the NMS with that value.
- Set a MIB variable—The SNMP agent begins this function in response to a message from the NMS. The SNMP agent changes the value of the MIB variable to the value requested by the NMS.

The SNMP agent also sends unsolicited trap messages to notify an NMS that a significant event has occurred on the agent. Examples of trap conditions include, but are not limited to, when a port goes up or down, when spanning-tree topology changes occur, and when authentication failures occur.

SNMP Community Strings

SNMP community strings authenticate access to MIB objects and function as embedded passwords. In order for the NMS to access the switch, the community string definitions on the NMS must match at least one of the three community string definitions on the switch.

A community string can have one of these attributes:

- Read-only (RO)—Gives read access to authorized management stations to all objects in the MIB except the community strings, but does not allow write access
- Read-write (RW)—Gives read and write access to authorized management stations to all objects in the MIB, but does not allow access to the community strings

- Read-write-all—Gives read and write access to authorized management stations to all objects in the MIB, including the community strings

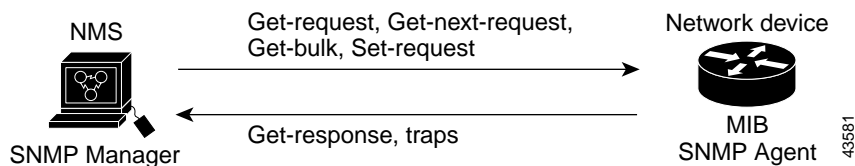
Note: When a cluster is created, the command switch manages the exchange of messages among member switches and the SNMP application. The Cluster Management software appends the member switch number (*@esN*, where *N* is the switch number) to the first configured RW and RO community strings on the command switch and propagates them to the member switches. For more information, see Chapter 5 “Clustering Switches.”

Using SNMP to Access MIB Variables

An example of an NMS is the CiscoWorks network management software. CiscoWorks 2000 software uses the switch MIB variables to set device variables and to poll devices on the network for specific information. The results of a poll can be displayed as a graph and analyzed to troubleshoot internetworking problems, increase network performance, verify the configuration of devices, monitor traffic loads, and more.

As shown in Figure 65, the SNMP agent gathers data from the MIB. The agent can send traps, or notification of certain events, to the SNMP manager, which receives and processes the traps. Traps alert the SNMP manager to a condition on the network such as improper user authentication, restarts, link status (up or down), MAC address tracking, and so forth. The SNMP agent also responds to MIB-related queries sent by the SNMP manager in *get-request*, *get-next-request*, and *set-request* format.

Figure 65. SNMP Network



For information on supported MIBs and how to access them, see Appendix A, “Supported MIBs.”

SNMP Notifications

SNMP allows the switch to send notifications to SNMP managers when particular events occur. SNMP notifications can be sent as traps or inform requests. In command syntax, unless there is an option in the command to select either traps or informs, the keyword *traps* refers to either traps or informs, or both. Use the **snmp-server host** command to specify whether to send SNMP notifications as traps or informs.

Traps are unreliable because the receiver does not send an acknowledgment when it receives a trap, and the sender cannot determine if the trap was received. When an SNMP manager receives an inform request, it acknowledges the message with an SNMP response protocol data unit (PDU). If the sender does not receive a response, the inform request can be sent again. Because they can be re-sent, informs are more likely than traps to reach their intended destination.

The characteristics that make informs more reliable than traps also consume more resources in the switch and in the network. Unlike a trap, which is discarded as soon as it is sent, an inform request is held in memory until a response is received or the request times out. Traps are sent only once, but an inform might be re-sent or retried several times. The retries increase traffic and contribute to a higher overhead on the

network. Therefore, traps and informs require a trade-off between reliability and resources. If it is important that the SNMP manager receive every notification, use inform requests. If traffic on the network or memory in the switch is a concern and notification is not required, use traps.

Note: SNMPv1 does not support informs.

Configuring SNMP

This section describes how to configure SNMP on your switch. It contains this configuration information:

- Default SNMP Configuration, on page 399
- SNMP Configuration Guidelines, on page 399
- Disabling the SNMP Agent, on page 400
- Configuring Community Strings, on page 400
- Configuring SNMP Groups and Users, on page 402
- Configuring SNMP Notifications, on page 404
- Setting the Agent Contact and Location Information, on page 407
- Limiting TFTP Servers Used Through SNMP, on page 407
- SNMP Examples, on page 408

Default SNMP Configuration

Table 80 shows the default SNMP configuration.

Table 80. Default SNMP Configuration

Feature	Default Setting
SNMP agent	Enabled
SNMP community strings	Read-Only: Public Read-Write: Private
SNMP trap receiver	None configured
SNMP traps	None enabled
SNMP version	If no version keyword is present, the default is version 1.
SNMPv3 authentication	If no keyword is entered, the default is the noauth (noAuthNoPriv) security level.
SNMP notification type	If no type is specified, all notifications are sent.

SNMP Configuration Guidelines

An SNMP *group* is a table that maps SNMP users to SNMP views. An SNMP *user* is a member of an SNMP group. An SNMP *host* is the recipient of an SNMP trap operation. An SNMP *engine ID* is a name for the local or remote SNMP engine.

When configuring SNMP, follow these guidelines:

- When configuring an SNMP group, do not specify a notify view. The **snmp-server host** global configuration command autogenerates a notify view for the user and then adds it to the group associated with that user. Modifying the group's notify view affects all users associated with that group. Refer to the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter*

System Command Reference for information about when you should configure notify views.

- To configure a remote user, specify the IP address or port number for the remote SNMP agent of the device where the user resides.
- Before you configure remote users for a particular agent, configure the SNMP engine ID, using the **snmp-server engineID** global configuration with the **remote** option. The remote agent's SNMP engine ID and user password are used to compute the authentication and privacy digests. If you do not configure the remote engine ID first, the configuration command fails.
- When configuring SNMP informs, you need to configure the SNMP engine ID for the remote agent in the SNMP database before you can send proxy requests or informs to it.
- Changing the value of the SNMP engine ID has important side effects. A user's password (entered on the command line) is converted to an MD5 or SHA security digest based on the password and the local engine ID. The command-line password is then destroyed, as required by RFC 2274. Because of this deletion, if the value of engineID changes, the security digests of SNMPv3 users become invalid, and you need to reconfigure SNMP users by using the **snmp-server user username** global configuration command. Similar restrictions require the reconfiguration of community strings when the engine ID changes.

Disabling the SNMP Agent

Beginning in privileged EXEC mode, follow these steps to disable the SNMP agent:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	no snmp-server	Disable the SNMP agent operation.
3.	end	Return to privileged EXEC mode.
4.	show running-config	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The **no snmp-server** global configuration command disables all running versions (version 1, version 2C, and version 3) on the device. No specific Cisco IOS command exists to enable SNMP. The first **snmp-server** global configuration command that you enter enables all versions of SNMP.

Configuring Community Strings

You use the SNMP community string to define the relationship between the SNMP manager and the agent. The community string acts like a password to permit access to the agent on the switch. Optionally, you can specify one or more of these characteristics associated with the string:

- An access list of IP addresses of the SNMP managers that are permitted to use the community string to gain access to the agent
- A MIB view, which defines the subset of all MIB objects accessible to the given community
- Read and write or read-only permission for the MIB objects accessible to the community

Beginning in privileged EXEC mode, follow these steps to configure a community string on the switch:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	snmp-server community <i>string</i> [view <i>view-name</i>] [ro rw] [<i>access-list-number</i>]	Configure the community string. <ul style="list-style-type: none"> For <i>string</i>, specify a string that acts like a password and permits access to the SNMP protocol. You can configure one or more community strings of any length. (Optional) For view, specify the view record accessible to the community. (Optional) Specify either read-only (ro) if you want authorized management stations to retrieve MIB objects, or specify read-write (rw) if you want authorized management stations to retrieve and modify MIB objects. By default, the community string permits read-only access to all objects. (Optional) For <i>access-list-number</i>, enter an IP standard access list numbered from 1 to 99 and 1300 to 1999.
3.	access-list <i>access-list-number</i> { deny / permit } <i>source</i> [<i>source-wildcard</i>]	(Optional) If you specified an IP standard access list number in Step 2, then create the list, repeating the command as many times as necessary. <ul style="list-style-type: none"> For <i>access-list-number</i>, enter the access list number specified in Step 2. The deny keyword denies access if the conditions are matched. The permit keyword permits access if the conditions are matched. For <i>source</i>, enter the IP address of the SNMP managers that are permitted to use the community string to gain access to the agent. (Optional) For <i>source-wildcard</i>, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Note: To disable access for an SNMP community, set the community string for that community to the null string (do not enter a value for the community string).

To remove a specific community string, use the **no snmp-server community** *string* global configuration command.

This example shows how to assign the string *comaccess* to SNMP, to allow read-only access, and to specify that IP access list 4 can use the community string to gain access to the switch SNMP agent:

```
Switch(config)# snmp-server community comaccess ro 4
```

Configuring SNMP Groups and Users

You can specify an identification name (*engineID*) for the local or remote SNMP server engine on the switch. You can configure an SNMP server group that maps SNMP users to SNMP views, and you can add new users to the SNMP group.

Beginning in privileged EXEC mode, follow these steps to configure SNMP on the switch:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	snmp-server engineID { local <i>engineid-string</i> remote <i>ip-address</i> [udp-port <i>port-number</i>] <i>engineid-string</i> }	Configure a name for either the local or remote copy of SNMP. <ul style="list-style-type: none"> The <i>engineid-string</i> is a 24-character ID string with the name of the copy of SNMP. You need not specify the entire 24-character engine ID if it contains trailing zeros. Specify only the portion of the engine ID up to the point where only zeros remain in the value. For example, to configure an engine ID of 123400000000000000000000, you can enter this: snmp-server engineID local 1234 If you select remote, specify the <i>ip-address</i> of the device that contains the remote copy of SNMP and the optional UDP port on the remote device. The default is 162.

Step	Command	Purpose
3.	snmp-server group <i>groupname</i> { v1 / v2c v3 [auth noauth priv]} [read <i>readview</i>] [write <i>writeview</i>] [notify <i>notifyview</i>] [access <i>access-list</i>]	<p>Configure a new SNMP group on the remote device.</p> <ul style="list-style-type: none"> • For <i>groupname</i>, specify the name of the group. • Specify a security model: <ul style="list-style-type: none"> — v1 is the least secure of the possible security models. — v2c is the second least secure model. It allows transmission of informs and integers twice the normal width. — v3, the most secure, requires you to select an authentication level: <ul style="list-style-type: none"> auth—Enables the Message Digest 5 (MD5) and the Secure Hash Algorithm (SHA) packet authentication. noauth —The noAuthNoPriv security level. This is the default if no keyword is specified. priv—Enables Data Encryption Standard (DES) packet encryption (also called <i>privacy</i>). <p>Note: The priv keyword is available only when the cryptographic software image is installed.</p> <ul style="list-style-type: none"> • (Optional) Enter read <i>readview</i> with a string (not to exceed 64 characters) that is the name of the view in which you can only view the contents of the agent. • (Optional) Enter write <i>writeview</i> with a string (not to exceed 64 characters) that is the name of the view in which you enter data and configure the contents of the agent. • (Optional) Enter notify <i>notifyview</i> with a string (not to exceed 64 characters) that is the name of the view in which you specify a notify, inform, or trap. • (Optional) Enter access <i>access-list</i> with a string (not to exceed 64 characters) that is the name of the access list.

Step	Command	Purpose
4.	snmp-server user <i>username</i> <i>groupname</i> [remote <i>host</i> [udp-port <i>port</i>]] { v1 / v2c v3 [auth { md5 sha } <i>auth-password</i>]} [encrypted] [access <i>access-list</i>]	Configure a new user to an SNMP group. <ul style="list-style-type: none"> The <i>username</i> is the name of the user on the host that connects to the agent. The <i>groupname</i> is the name of the group to which the user is associated. (Optional) Enter remote to specify a remote SNMP entity to which the user belongs and the hostname or IP address of that entity with the optional UDP port number. The default is 162. Enter the SNMP version number (v1, v2c, or v3). If you enter v3, you have these additional options: <ul style="list-style-type: none"> auth is an authentication level setting session, which can be either the HMAC-MD5-96 or the HMAC-SHA-96 authentication level, and requires a password string (not to exceed 64 characters). encrypted specifies that the password appears in encrypted format. (Optional) Enter access <i>access-list</i> with a string (not to exceed 64 characters) that is the name of the access list.
5.	end	Return to privileged EXEC mode.
6.	show running-config	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Configuring SNMP Notifications

A trap manager is a management station that receives and processes traps. Traps are system alerts that the switch generates when certain events occur. By default, no trap manager is defined, and no traps are sent. Switches running this software release can have an unlimited number of trap managers.

Note: Many commands use the word *traps* in the command syntax. Unless there is an option in the command to select either traps or informs, the keyword *traps* refers to either traps, informs, or both. Use the **snmp-server host** command to specify whether to send SNMP notifications as traps or informs.

Table 81 describes the supported switch traps (notification types). You can enable any or all of these traps and configure a trap manager to receive them.

Table 81. Switch Notification Types

Notification Type Keyword	Description
bridge	Generates STP bridge MIB traps.
c2900	Generates a trap for switch-specific notifications.
cluster	Generates a trap when the cluster configuration changes.
config	Generates a trap for SNMP configuration changes.
copy-config	Generates a trap for SNMP copy configuration changes.
entity	Generates a trap for SNMP entity changes.
envmon	Generates environmental monitor traps. You can enable any or all of these environmental traps: fan, shutdown, supply, temperature.

Table 81. Switch Notification Types (continued)

Notification Type Keyword	Description
flash	Generates SNMP flash notifications.
hsrp	Generates a trap for Hot Standby Router Protocol (HSRP) changes.
mac-notification	Generates a trap for MAC address notifications.
rtr	Generates a trap for the SNMP Response Time Reporter (RTR).
snmp	Generates a trap for SNMP-type notifications.
stpx	Generates SNMP STP Extended MIB traps.
syslog	Generates SNMP syslog traps.
tty	Sends Cisco enterprise-specific notifications when a Transmission Control Protocol (TCP) connection closes.
udp-port	Sends notification of the User Datagram Protocol (UDP) port number of the host.
vlan-membership	Generates a trap for SNMP VLAN membership changes.
vlancreate	Generates SNMP VLAN-created traps.
vlandelete	Generates SNMP VLAN-deleted traps.
vtp	Generates a trap for VLAN Trunking Protocol (VTP) changes.

Some notification types cannot be controlled with the **snmp-server enable** global configuration command, for example, **tty** and **udp-port**. These notification types are always enabled. You can use the **snmp-server host** global configuration command to a specific host to receive the notification types listed in Table 81.

Beginning in privileged EXEC mode, follow these steps to configure the switch to send traps or informs to a host:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	snmp-server engineID remote <i>ip-address engineid-string</i>	Specify the engine ID for the remote host.
3.	snmp-server user <i>username</i> <i>groupname remote host</i> [udp-port <i>port</i>] { v1 / v2c v3 [auth { md5 sha } <i>auth-password</i>]} [encrypted] [access <i>access-list</i>]	Configure an SNMP user to be associated with the remote host created in Step 2. Note: You cannot configure a remote user for an address without first configuring the engine ID for the remote host. If you try to configure the user before configuring the remote engine ID, you receive an error message, and the command is not executed.
4.	snmp-server group [<i>groupname</i> { v1 / v2c v3 [auth noauth]}] [read <i>readview</i>] [write <i>writeview</i>] [notify <i>notifyview</i>] [access <i>access-list</i>]	Configure an SNMP group.

Step	Command	Purpose
5.	snmp-server host <i>host-addr</i> [traps informs] [version {1 / 2c 3 [auth noauth priv]] <i>community-string</i> [udp-port <i>port</i>] [notification-type]	Specify the recipient of an SNMP trap operation. <ul style="list-style-type: none"> For <i>host-addr</i>, specify the name or Internet address of the host (the targeted recipient). (Optional) Enter traps (the default) to send SNMP traps to the host. (Optional) Enter informs to send SNMP informs to the host. (Optional) Specify the SNMP version (1, 2c, or 3). SNMPv1 is not available with informs. (Optional) For version 3, select authentication level auth, noauth, or priv. <p>Note: The priv keyword is available only when the cryptographic software image is installed.</p> <ul style="list-style-type: none"> For <i>community-string</i>, enter the password-like community string sent with the notification operation. (Optional) For udp-port port, enter the UDP port on the remote device. (Optional) For <i>notification-type</i>, use the keywords listed in Table 81 on page 404. If no type is specified, all notifications are sent.
6.	snmp-server enable traps <i>notification-types</i>	Enable the switch to send traps or informs and specify the type of notifications to be sent. For a list of notification types, see Table 81 on page 404, or enter this: snmp-server enable traps ? To enable multiple types of traps, you must enter a separate snmp-server enable traps command for each trap type.
7.	snmp-server trap-source <i>interface-id</i>	(Optional) Specify the source interface, which provides the IP address for the trap message. This command also sets the source IP address for informs.
8.	snmp-server queue-length <i>length</i>	(Optional) Establish the message queue length for each trap host. The range is 1 to 1000; the default is 10.
9.	snmp-server trap-timeout <i>seconds</i>	(Optional) Define how often to resend trap messages. The range is 1 to 1000; the default is 30 seconds.
10.	end	Return to privileged EXEC mode.
11.	show running-config	Verify your entries.
12.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The **snmp-server host** command specifies which hosts receive the notifications. The **snmp-server enable trap** command globally enables the mechanism for the specified notification (for traps and informs). To enable a host to receive an inform, you must configure an **snmp-server host informs** command for the host and globally enable informs by using the **snmp-server enable traps** command.

To remove the specified host from receiving traps, use the **no snmp-server host host** global configuration command. The **no snmp-server host** command with no keywords disables traps, but not informs, to the host. To disable informs, use the **no snmp-server host informs** global configuration command. To disable a specific trap type, use the **no snmp-server enable traps notification-types** global configuration command.

Setting the Agent Contact and Location Information

Beginning in privileged EXEC mode, follow these steps to set the system contact and location of the SNMP agent so that these descriptions can be accessed through the configuration file:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	snmp-server contact <i>text</i>	Set the system contact string. For example: snmp-server contact Dial System Operator at beeper 21555.
3.	snmp-server location <i>text</i>	Set the system location string. For example: snmp-server location Building 3/Room 222
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Limiting TFTP Servers Used Through SNMP

Beginning in privileged EXEC mode, follow these steps to limit the TFTP servers used for saving and loading configuration files through SNMP to the servers specified in an access list:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	snmp-server tftp-server-list <i>access-list-number</i>	Limit TFTP servers used for configuration file copies through SNMP to the servers in the access list. For <i>access-list-number</i> , enter an IP standard access list numbered from 1 to 99 and 1300 to 1999.
3.	access-list <i>access-list-number</i> { deny / permit } <i>source</i> [<i>source-wildcard</i>]	Create a standard access list, repeating the command as many times as necessary. <ul style="list-style-type: none"> For <i>access-list-number</i>, enter the access list number specified in Step 2. The deny keyword denies access if the conditions are matched. The permit keyword permits access if the conditions are matched. For <i>source</i>, enter the IP address of the TFTP servers that can access the switch. (Optional) For <i>source-wildcard</i>, enter the wildcard bits, in dotted decimal notation, to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
4.	end	Return to privileged EXEC mode.

Step	Command	Purpose
5.	show running-config	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

SNMP Examples

This example shows how to enable all versions of SNMP. The configuration permits any SNMP manager to access all objects with read-only permissions using the community string *public*. This configuration does not cause the switch to send any traps.

```
Switch(config)# snmp-server community public
```

This example shows how to permit any SNMP manager to access all objects with read-only permission using the community string *public*. The switch also sends VTP traps to the hosts 192.180.1.111 and 192.180.1.33 using SNMPv1 and to the host 192.180.1.27 using SNMPv2C. The community string *public* is sent with the traps.

```
Switch(config)# snmp-server community public
Switch(config)# snmp-server enable traps vtp
Switch(config)# snmp-server host 192.180.1.27 version 2c public
Switch(config)# snmp-server host 192.180.1.111 version 1 public
Switch(config)# snmp-server host 192.180.1.33 public
```

This example shows how to allow read-only access for all objects to members of access list 4 that use the *comaccess* community string. No other SNMP managers have access to any objects. SNMP Authentication Failure traps are sent by SNMPv2C to the host *ibm.com* using the community string *public*.

```
Switch(config)# snmp-server community comaccess ro 4
Switch(config)# snmp-server enable traps snmp authentication
Switch(config)# snmp-server host ibm.com version 2c public
```

This example shows how to send Entity MIB traps to the host *ibm.com*. The community string is restricted. The first line enables the switch to send Entity MIB traps in addition to any traps previously enabled. The second line specifies the destination of these traps and overwrites any previous **snmp-server host** commands for the host *ibm.com*.

```
Switch(config)# snmp-server enable traps entity
Switch(config)# snmp-server host ibm.com restricted entity
```

This example shows how to enable the switch to send all traps to the host *myhost.ibm.com* using the community string *public*:

```
Switch(config)# snmp-server enable traps
Switch(config)# snmp-server host myhost.ibm.com public
```

Displaying SNMP Status

To display SNMP input and output statistics, including the number of illegal community string entries, errors, and requested variables, use the **show snmp** privileged EXEC command. You can also use the other privileged EXEC commands in Table 82 to display SNMP information. For information about the fields in the output displays, refer to the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference*.

Table 82. Commands for Displaying SNMP Information

Feature	Default Setting
show snmp	Displays SNMP statistics.
show snmp engineID [local remote]	Displays information on the local SNMP engine and all remote engines that have been configured on the device.
show snmp group	Displays information on each SNMP group on the network.
show snmp user	Displays information on each SNMP user name in the SNMP users table.

Chapter 24. Configuring QoS

This chapter describes how to configure quality of service (QoS) by using automatic-QoS (auto-QoS) commands or by using standard QoS commands. With QoS, you can give preferential treatment to certain types of traffic at the expense of others. Without QoS, the IGESM offers best-effort service to each packet, regardless of the packet contents or size. It sends the packets without any assurance of reliability, delay bounds, or throughput.

Note: For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

QoS can be configured either by using the Cluster Management Suite (CMS) or through the command-line interface (CLI). Refer to the CMS online help for configuration procedures through CMS. For information about accessing and using CMS, see Chapter 3 “Getting Started with CMS.”

You can also use these wizards to configure QoS:

- Priority data wizard—Lets you assign priority levels to data applications based on their TCP or UDP ports. It has a standard list of applications, and you select the ones that you want to prioritize, the priority levels, and the interfaces where the prioritization occurs. Refer to the priority data wizard online help for procedures about using this wizard.
- Video wizard—Gives traffic that originates from specified video servers a higher priority than the priority of data traffic. The wizard assumes that the video servers are connected to a single device in the cluster. Refer to the video wizard online help for procedures about using this wizard.

This chapter consists of these sections:

- Understanding QoS, on page 411
- Configuring Auto-QoS, on page 417
- Displaying Auto-QoS Information, on page 421
- Configuring Standard QoS, on page 421
- Displaying Standard QoS Information, on page 440
- Standard QoS Configuration Examples, on page 440

Understanding QoS

This section describes how QoS is implemented on the switch. Typically, networks operate on a best-effort delivery basis, which means that all traffic has equal priority and an equal chance of being delivered in a timely manner. When congestion occurs, all traffic has an equal chance of being dropped.

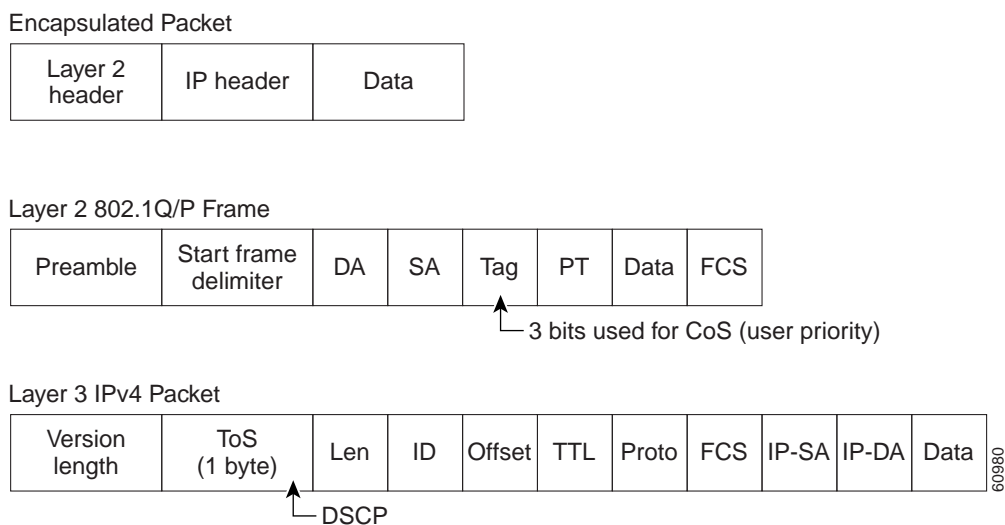
When you configure the QoS feature, you can select specific network traffic, prioritize it according to its relative importance, and use congestion-management and congestion-avoidance techniques to give preferential treatment. Implementing QoS in your network makes network performance more predictable and bandwidth utilization more effective.

The QoS implementation is based on the DiffServ architecture, an emerging standard from the Internet Engineering Task Force (IETF). This architecture specifies that each packet is classified upon entry into the network. The classification is carried in the IP packet header, using 6 bits from the deprecated IP type-of-service (ToS) field to carry the classification (*class*) information.

Classification can also be carried in the Layer 2 frame. These special bits in the Layer 2 frame or a Layer 3 packet are described here and shown in Figure 66:

- **Prioritization values in Layer 2 frames**
 Layer 2 802.1Q frame headers have a 2-byte Tag Control Information field that carries the class of service (CoS) value in the three most-significant bits, which are called the User Priority bits. On interfaces configured as Layer 2 802.1Q trunks, all traffic is in 802.1Q frames except for traffic in the native VLAN.
 Other frame types cannot carry Layer 2 CoS values.
 Layer 2 CoS values range from 0 for low priority to 7 for high priority.
- **Prioritization bits in Layer 3 packets**
 Layer 3 IP packets can carry a Differentiated Services Code Point (DSCP) value. The supported DSCP values are 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56.

Figure 66. QoS Classification Layers in Frames and Packets



All switches and routers that access the Internet rely on the class information to give the same forwarding treatment to packets with the same class information and different treatment to packets with different class information. The class information in the packet can be assigned by end hosts or by switches or routers along the way, based on a configured policy, detailed examination of the packet, or both. Detailed examination of the packet is expected to happen closer to the edge of the network so that the core switches and routers are not overloaded.

Switches and routers along the path can use the class information to limit the amount of resources allocated per traffic class. The behavior of an individual device when handling traffic in the DiffServ architecture is called per-hop behavior. If all devices along a path have a consistent per-hop behavior, you can construct an end-to-end QoS solution.

Implementing QoS in your network can be a simple or complex task and depends on the QoS features offered by your internetworking devices, the traffic types and patterns in your network, and the granularity of control that you need over incoming and outgoing traffic.

Basic QoS Model

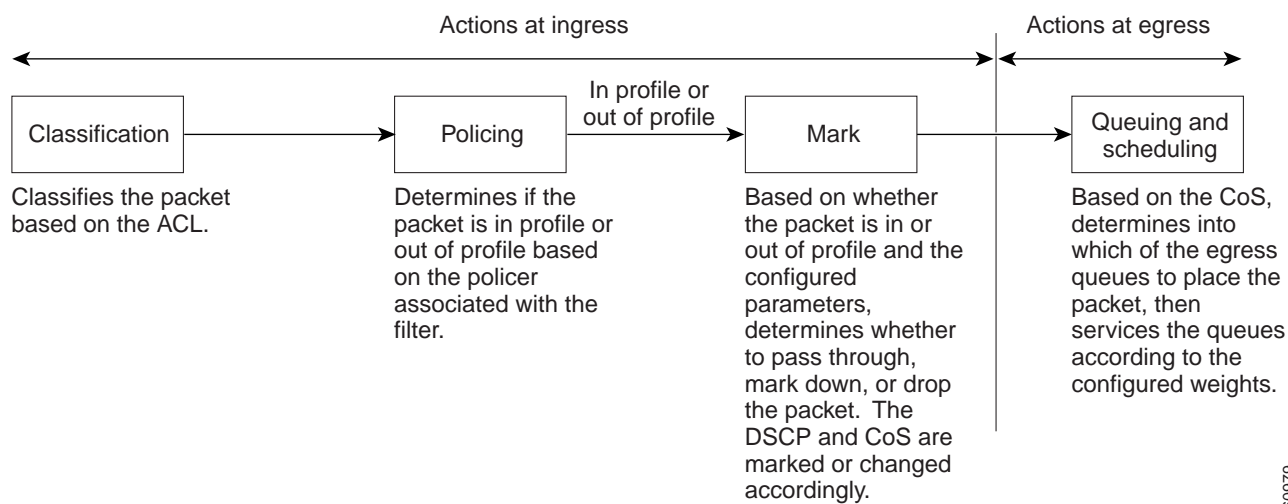
Figure 67 shows the basic QoS model. Actions at the ingress interface include classifying traffic, policing, and marking:

- Classifying distinguishes one kind of traffic from another. For more information, see the “Classification” section on page 413.
- Policing determines whether a packet is in or out of profile according to the configured policer, and the policer limits the bandwidth consumed by a flow of traffic. The result of this determination is passed to the marker. For more information, see the “Policing and Marking” section on page 415.
- Marking evaluates the policer and configuration information for the action to be taken when a packet is out of profile and decides what to do with the packet (pass through a packet without modification, mark down the DSCP value in the packet, or drop the packet). For more information, see the “Policing and Marking” section on page 415.

Actions at the egress interface include queueing and scheduling:

- Queueing evaluates the CoS value and determines which of the four egress queues in which to place the packet.
- Scheduling services the four egress queues based on their configured weighted round robin (WRR) weights.

Figure 67. Basic QoS Model.



Classification

Classification is the process of distinguishing one kind of traffic from another by examining the fields in the packet.

Classification occurs only on a physical interface basis. No support exists for classifying packets at the VLAN level.

You specify which fields in the frame or packet that you want to use to classify incoming traffic.

For non-IP traffic, you have these classification options:

- Use the port default. If the frame does not contain a CoS value, the switch assigns the default port CoS value to the incoming frame.

- Trust the CoS value in the incoming frame (configure the port to trust CoS). Layer 2 802.1Q frame headers carry the CoS value in the three most-significant bits of the Tag Control Information field. CoS values range from 0 for low priority to 7 for high priority.

The trust DSCP configuration is meaningless for non-IP traffic. If you configure a port with this option and non-IP traffic is received, the switch assigns the default port CoS value and classifies traffic based on the CoS value.

For IP traffic, you have these classification options:

- Trust the IP DSCP in the incoming packet (configure the port to trust DSCP). The switch assigns the same DSCP to the packet for internal use. The IETF defines the 6 most-significant bits of the 1-byte ToS field as the DSCP. The priority represented by a particular DSCP value is configurable. The supported DSCP values are 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56.
- Trust the CoS value (if present) in the incoming packet. The switch generates the DSCP by using the CoS-to-DSCP map.

Note: An interface can be configured to trust either CoS or DSCP, but not both at the same time.

Classification Based on QoS ACLs

You can use IP standard, IP extended, and Layer 2 MAC access control lists (ACLs) to define a group of packets with the same characteristics (*class*). In the QoS context, the permit and deny actions in the access control entries (ACEs) have different meanings than with security ACLs:

- If a match with a permit action is encountered (first-match principle), the specified QoS-related action is taken.
- If no match with a permit action is encountered and all the ACEs have been examined, no QoS processing occurs on the packet.
- If multiple ACLs are configured on an interface, the packet matches the first ACL with a permit action, and QoS processing begins.
- Configuration of a deny action is not supported in QoS ACLs on the switch.
- System-defined masks are allowed in class maps with these restrictions:
 - A combination of system-defined and user-defined masks cannot be used in the multiple class maps that are a part of a policy map.
 - System-defined masks that are a part of a policy map must all use the same type of system mask. For example, a policy map cannot have a class map that uses the **permit tcp any any** ACE and another that uses the **permit ip any any** ACE.
 - A policy map can contain multiple class maps that all use the same user-defined mask or the same system-defined mask.

Note: For more information about system-defined masks, see the “Understanding Access Control Parameters” section on page 373.

For more information about ACL restrictions, see the “Configuring ACLs” section on page 375.

After a traffic class has been defined with the ACL, you can attach a policy to it. A policy might contain multiple classes with actions specified for each one of them. A policy might include commands to classify the class as a particular aggregate (for example, assign a DSCP) or rate-limit the class. This policy is then attached to a particular port on which it becomes effective.

You implement IP ACLs to classify IP traffic by using the **access-list** global configuration command; you implement Layer 2 MAC ACLs to classify Layer 2 traffic by using the **mac access-list extended** global configuration command.

Classification Based on Class Maps and Policy Maps

A class map is a mechanism that you use to isolate and name a specific traffic flow (or class) from all other traffic. The class map defines the criteria used to match against a specific traffic flow to further classify it; the criteria can include matching the access group defined by the ACL. If you have more than one type of traffic that you want to classify, you can create another class map and use a different name. After a packet is matched against the class-map criteria, you further classify it through the use of a policy map.

A policy map specifies which traffic class to act on. Actions can include setting a specific DSCP value in the traffic class or specifying the traffic bandwidth limitations and the action to take when the traffic is out of profile. Before a policy map can be effective, you must attach it to an interface.

You create a class map by using the **class-map** global configuration command or the **class** policy-map configuration command. You should use the **class-map** global configuration command when the map is shared among many ports. When you enter the **class-map** global configuration command, the switch enters the class-map configuration mode. In this mode, you define the match criterion for the traffic by using the **match** class-map configuration command.

You create and name a policy map by using the **policy-map** global configuration command. When you enter this command, the switch enters the policy-map configuration mode. In this mode, you specify the actions to take on a specific traffic class by using the **class** policy-map configuration or **set** policy-map class configuration command. To make the policy map effective, you attach it to an interface by using the **service-policy** interface configuration command.

The policy map can also contain commands that define the policer, the bandwidth limitations of the traffic, and the action to take if the limits are exceeded. For more information, see the “Policing and Marking” section on page 415.

A policy map also has these characteristics:

- A policy map can contain multiple class statements.
- A separate policy-map class can exist for each type of traffic received through an interface.
- A policy-map configuration state supersedes any actions due to an interface trust state.

For configuration information, see the “Configuring a QoS Policy” section on page 428.

Policing and Marking

Policing involves creating a policer that specifies the bandwidth limits for the traffic. Packets that exceed the limits are *out of profile* or *nonconforming*. Each policer specifies the action to take for packets that are in or out of profile. These actions, carried out by the marker, include dropping the packet or marking down the packet with a new user-defined value.

You can create an individual policer. QoS applies the bandwidth limits specified in the policer separately to each matched traffic class. You configure this type of policer within a policy map by using the **policy-map** configuration command.

When configuring policing and policers, keep these items in mind:

- By default, no policers are configured.
- Policers can only be configured on a physical port. There is no support for policing at a VLAN level.
- Only one policer can be applied to a packet in the input direction.
- Only the average rate and committed burst parameters are configurable.
- Policing occurs on the ingress interfaces:
 - 60 policers are supported on ingress Gigabit-capable Ethernet ports.
 - Granularity for the average burst rate is 8 Mbps for Gigabit Ethernet ports.
- On an interface configured for QoS, all traffic received through the interface is classified, policed, and marked according to the policy map attached to the interface. On a trunk interface configured for QoS, traffic in *all* VLANs received through the interface is classified, policed, and marked according to the policy map attached to the interface.

Note: You cannot configure policers on the egress interfaces.

Mapping Tables

During classification, QoS uses a configurable CoS-to-DSCP map to derive an internal DSCP value from the received CoS value. This DSCP value represents the priority of the traffic.

Before the traffic reaches the scheduling stage, QoS uses the configurable DSCP-to-CoS map to derive a CoS value from the internal DSCP value. The CoS value is used to select one of the four egress queues.

The CoS-to-DSCP and DSCP-to-CoS maps have default values that might or might not be appropriate for your network.

For configuration information, see the “Configuring CoS Maps” section on page 436.

Queueing and Scheduling

The switch gives QoS-based 802.1p CoS values. QoS uses classification and scheduling to send network traffic from the switch in a predictable manner. QoS classifies frames by assigning priority-indexed CoS values to them and gives preference to higher-priority traffic such as telephone calls.

How Class of Service Works

Before you set up 802.1p CoS on a CIGESM switch that operates with the Catalyst 6000 family of switches, refer to the Catalyst 6000 documentation. There are differences in the 802.1p implementation that you should understand to ensure compatibility.

Port Priority

Frames received from users in the administratively-defined VLANs are classified or *tagged* for transmission to other devices. Based on rules that you define, a unique identifier (the tag) is inserted in each frame header before it is forwarded. The tag is examined and understood by each device before any broadcasts or transmissions to other switches, routers, or end stations. When the frame reaches the last switch or router, the tag is removed before the frame is sent to the target end station. VLANs

that are assigned on trunk or access ports without identification or a tag are called *native* or *untagged* frames.

For IEEE 802.1Q frames with tag information, the priority value from the header frame is used. For native frames, the default priority of the input port is used.

Port Scheduling

Each port on the switch has a single receive queue buffer (the *ingress* port) for incoming traffic. When an untagged frame arrives, it is assigned the value of the port as its port default priority. You assign this value by using the CLI or CMS. A tagged frame continues to use its assigned CoS value when it passes through the ingress port.

CoS configures each transmit port (the *egress* port) with a normal-priority transmit queue and a high-priority transmit queue, depending on the frame tag or the port information. Frames in the normal-priority queue are forwarded only after frames in the high-priority queue are forwarded.

The switch (802.1p user priority) has four priority queues. The frames are forwarded to appropriate queues based on the priority-to-queue mapping that you defined.

Egress CoS Queues

The switch supports four CoS queues for each egress port. For each queue, you can specify these types of scheduling:

- Strict priority scheduling

Strict priority scheduling is based on the priority of queues. Packets in the high-priority queue always transmit first, and packets in the low-priority queue do not transmit until all the high-priority queues become empty.

The default scheduling method is strict priority.
- Weighted round-robin (WRR) scheduling

WRR scheduling requires you to specify a number that indicates the importance (weight) of the queue relative to the other CoS queues. WRR scheduling prevents the low-priority queues from being completely neglected during periods of high-priority traffic. The WRR scheduler transmits some packets from each queue in turn. The number of packets it sends corresponds to the relative importance of the queue. For example, if one queue has a weight of 3 and another has a weight of 4, three packets are sent from the first queue for every four that are sent from the second queue. By using this scheduling, low-priority queues have the opportunity to send packets even though the high-priority queues are not empty.
- Strict priority and WRR scheduling

Strict priority and WRR scheduling, also referred to as strict priority queueing, uses one of the egress queues as an expedite queue (queue 4). The remaining queues participate in WRR. When the expedite queue is configured, it is a priority queue and is serviced until it is empty before the other queues are serviced by WRR scheduling.

You can enable the egress expedite queue and assign WRR weights to the other queues by using the **wrr-queue bandwidth weight1 weight2 weight3 0** global configuration command.

Configuring Auto-QoS

You can use the auto-QoS feature to simplify the deployment of existing QoS features. Auto-QoS makes assumptions about the network design, and as a result, the switch

can prioritize different traffic flows and appropriately use the egress queues instead of using the default QoS behavior (the switch offers best-effort service to each packet regardless of the packet contents or size and sends it from a single queue).

When you enable auto-QoS, it automatically classifies traffic based on the traffic type and ingress packet label. The switch uses the resulting classification to choose the appropriate egress queue.

You use auto-QoS commands to identify ports connected to Cisco IP Phones and to identify ports that receive trusted voice over IP (VoIP) traffic through an uplink. Auto-QoS then performs these functions:

- Detects the presence or absence of IP phones
- Configures QoS classification
- Configures egress queues

These sections describe how to configure auto-QoS on your switch:

- Generated Auto-QoS Configuration, on page 418
- Effects of Auto-QoS on the Configuration, on page 420
- Configuration Guidelines, on page 420
- Enabling Auto-QoS for VoIP, on page 420

Generated Auto-QoS Configuration

When auto-QoS is enabled, it uses the ingress packet label to classify traffic and to configure the egress queues as described in Table 83.

Table 83. Traffic Types, Ingress Packet Labels, Assigned Packet Labels, and Egress Queues

	VoIP Data Traffic Only From Cisco IP Phones	VoIP Control Traffic Only From Cisco IP Phones	Routing Protocol Traffic	STP BPDU¹ Traffic	All Other Traffic
Ingress DSCP	46	26	–	–	–
Ingress CoS	5	3	6	7	–
Assigned DSCP	46	26	48	56	0
Assigned CoS	5	3	6	7	0
CoS-to-Queue Map	5	3, 6, 7			0, 1, 2, 4
Egress Queue	Expedite queue	80% WRR			20% WRR

1. BPDU = bridge protocol data unit

Table 84 lists the generated auto-QoS configuration for the egress queues.

Table 84. Auto-QoS Configuration for the Egress Queues

Egress Queue	Queue Number	CoS-to-Queue Map	Queue Weight
Expedite	4	5	–
80% WRR	3	3, 6, 7	80%
20% WRR	1	0, 1, 2, 4	20%

When you enable the auto-QoS feature on the first interface, these automatic actions occur:

- When you enter the **auto qos voip trust** interface configuration command, the ingress classification on the interface is set to trust the QoS label received in the packet, and the egress queues on the interface are reconfigured (see Table 84).
- When you enter the **auto qos voip cisco-phone** interface configuration command, the trusted boundary feature is enabled. It uses the Cisco Discovery Protocol (CDP) to detect the presence or absence of a Cisco IP Phone. When a Cisco IP Phone is detected, the ingress classification on the interface is set to trust the QoS label received in the packet. When a Cisco IP Phone is absent, the ingress classification is set to not trust the QoS label in the packet. The egress queues on the interface are reconfigured (see Table 84).

For information about the trusted boundary feature, see the “Configuring Trusted Boundary” section on page 426.

- The switch automatically assigns egress queue usage as shown in Table 84.

When you enable auto-QoS by using the **auto qos voip cisco-phone** or the **auto qos voip trust** interface configuration command, the switch automatically generates a QoS configuration based on the traffic type and ingress packet label and applies the commands listed in Table 85 to the interface.

Table 85. Generated Auto-QoS Configuration Command Equivalents

Description	Automatically Generated QoS Command Equivalent
The switch automatically enables standard QoS and configures the CoS-to-DSCP map (maps CoS values in incoming packets to a DSCP value) as shown in Table 83 on page 418.	Switch(config)# mls qos map cos-dscp 0 8 16 26 32 46 48 56
The switch automatically sets the ingress classification on the interface to trust the CoS value received in the packet.	Switch(config-if)# mls qos trust cos
If you entered the auto qos voip cisco-phone command, the switch automatically enables the trusted boundary feature, which uses the CDP to detect the presence or absence of a Cisco IP Phone.	Switch(config-if)# mls qos trust device cisco-phone
<p>The switch automatically assigns egress queue usage (as shown in Table 84 on page 418) on this interface.</p> <p>The switch enables the egress expedite queue and assigns WRR weights to queues 1 and 3. (The lowest value for a WRR queue is 1. When the WRR weight of a queue is set to 0, this queue becomes an expedite queue.)</p> <p>The switch configures the CoS-to-egress-queue map:</p> <ul style="list-style-type: none"> • CoS values 0, 1, 2, and 4 select queue 1. • CoS values 3, 6, and 7 select queue 3. • CoS value 5 selects queue 4 (expedite queue). <p>Because the expedite queue (queue 4) contains the VoIP data traffic, the queue is serviced until empty.</p>	<p>Switch(config)# wrr-queue bandwidth 20 1 80 0</p> <p>Switch(config)# wrr-queue cos-map 1 0 1 2 4</p> <p>Switch(config)# wrr-queue cos-map 3 3 6 7</p> <p>Switch(config)# wrr-queue cos-map 4 5</p>

Effects of Auto-QoS on the Configuration

When auto-QoS is enabled, the **auto qos voip** interface configuration command and the generated configuration are added to the running configuration.

Configuration Guidelines

Before configuring auto-QoS, you should be aware of this information:

- In this release, auto-QoS configures the switch only for VoIP with Cisco IP Phones.
- To take advantage of the auto-QoS defaults, do not configure any standard QoS commands before entering the auto-QoS commands. If necessary, you can fine-tune the QoS configuration, but we recommend that you do so only after the auto-QoS configuration is completed.
- You can enable auto-QoS on static, dynamic-access, voice VLAN access, and trunk ports.
- By default, the CDP is enabled on all interfaces. For auto-QoS to function properly, do not disable the CDP.
- Policing is not enabled in auto-QoS. You can manually enable policing, as described in the “Configuring a QoS Policy” section on page 428.

Enabling Auto-QoS for VoIP

Beginning in privileged EXEC mode, follow these steps to enable auto-QoS for VoIP within a QoS domain:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface that is connected to a Cisco IP Phone. You also can specify the uplink interface that is connected to another switch or router in the interior of the network.
3.	auto qos voip {cisco-phone trust}	Enable auto-QoS. The keywords have these meanings: <ul style="list-style-type: none">• cisco-phone—If the interface is connected to a Cisco IP Phone, the QoS labels of incoming packets are trusted only when the IP phone is detected.• trust—The uplink interface is connected to a trusted switch or router, and the VoIP classification in the ingress packet is trusted.
4.	end	Return to privileged EXEC mode.
5.	show auto qos interface <i>interface-id</i>	Verify your entries. This command displays the auto-QoS configuration that was initially applied; it does not display any user changes to the configuration that might be in effect.

To disable auto-QoS on the switch and return to the default port trust state set (untrusted), follow these steps:

1. Use the **no auto qos voip** interface configuration command on all interfaces on which auto-QoS is enabled. To disable auto-QoS on multiple interfaces at the same time, you can use the **interface range** global configuration command.
2. After disabling auto-QoS on all interfaces on which auto-QoS was enabled, return the egress queues and CoS-to-DSCP map to the default settings by using these global configuration commands:
 - **no wrr-queue bandwidth**
 - **no wrr-queue cos-map**
 - **no mls qos map cos-dscp**

To display the QoS commands that are automatically generated when auto-QoS is enabled or disabled, enter the **debug autoqos** privileged EXEC command before enabling auto-QoS. For more information, see the “Using the debug autoqos Command” section on page 462.

This example shows how to enable auto-QoS and to trust the QoS labels in incoming packets when the device connected to Gigabit Ethernet interface 0/17 is detected as a Cisco IP Phone:

```
Switch(config)# interface gigabitethernet0/17  
Switch(config-if)# auto qos voip cisco-phone
```

This example shows how to enable auto-QoS and to trust the QoS labels in incoming packets when the switch or router connected to Gigabit Ethernet interface 0/17 is a trusted device:

```
Switch(config)# interface gigabitethernet0/17  
Switch(config-if)# auto qos voip trust
```

Displaying Auto-QoS Information

To display the initial auto-QoS configuration, use the **show auto qos [interface *interface-id*]** privileged EXEC command. To display any user changes to that configuration, use the **show running-config** privileged EXEC command. You can compare the **show auto qos** and the **show running-config** command output to identify the user-defined QoS settings.

To display information about the QoS configuration that might be affected by auto-QoS, use one of these commands:

- **show mls qos**
- **show mls qos map cos-dscp**
- **show wrr-queue bandwidth**
- **show wrr-queue cos-map**

For more information about these commands, refer to the command reference for this release.

Configuring Standard QoS

Before configuring standard QoS, you must have a thorough understanding of these items:

- The types of applications used and the traffic patterns on your network.
- Traffic characteristics and needs of your network. Is the traffic bursty? Do you need to reserve bandwidth for voice and video streams?
- Bandwidth requirements and speed of the network.
- Location of congestion points in the network.

This section describes how to configure standard QoS on your switch:

- Default Standard QoS Configuration, on page 422
- Configuration Guidelines, on page 422
- Configuring Classification Using Port Trust States, on page 423
- Configuring a QoS Policy, on page 428
- Configuring CoS Maps, on page 436
- Configuring the Egress Queues, on page 438

Default Standard QoS Configuration

This is the default standard QoS configuration:

- The default port CoS value is 0.
- The CoS value of 0 is assigned to all incoming packets.
- The default port trust state is untrusted.
- No policy maps are configured.
- No policers are configured.
- The default CoS-to-DSCP map is shown in Table 88.
- The default DSCP-to-CoS map is shown in Table 89.
- The default scheduling method for the egress queues is strict priority.
- For default CoS and WRR values, see the “Configuring the Egress Queues” section on page 438.

Configuration Guidelines

Before beginning the QoS configuration, you should be aware of this information:

- You must disable the IEEE 802.3x flowcontrol on all ports before enabling QoS on the switch. To disable it, use the **flowcontrol receive off** and **flowcontrol send off** interface configuration commands.
- If you have EtherChannel ports configured on your switch, you must configure QoS classification, policing, mapping, and queueing on the individual physical ports that comprise the EtherChannel. You must decide whether the QoS configuration should match on all ports in the EtherChannel.
- It is not possible to match IP fragments against configured IP extended ACLs to enforce QoS. IP fragments are sent as best-effort traffic. IP fragments are denoted by fields in the IP header.
- All ingress QoS processing actions apply to control traffic (such as spanning-tree bridge protocol data units [BPDUs] and routing update packets) that the switch receives.
- Only an ACL that is created for physical interfaces can be attached to a class map.

- Only one ACL per class map and only one **match** command per class map are supported. The ACL can have multiple access control entries, which are commands that match fields against the contents of the packet.
- Policy maps with ACL classification in the egress direction are not supported and cannot be attached to an interface by using the **service-policy input** *policy-map-name* interface configuration command.
- In a policy map, the class named *class-default* is not supported. The switch does not filter traffic based on the policy map defined by the **class class-default** policy-map configuration command.
- For more information about guidelines for configuring ACLs, see the “Classification Based on QoS ACLs” section on page 414.
- For information about applying ACLs to physical interfaces, see the “Guidelines for Applying ACLs to Physical Interfaces” section on page 375.
- If a policy map with a system-defined mask and a security ACL with a user-defined mask are configured on an interface, the switch might ignore the actions specified by the policy map and perform only the actions specified by the ACL. For information about masks, see the “Understanding Access Control Parameters” section on page 373.
- If a policy map with a user-defined mask and a security ACL with a user-defined mask are configured on an interface, the switch takes one of the actions as described in Table 86. For information about masks, see the “Understanding Access Control Parameters” section on page 373.

Table 86. Interaction Between Policy Maps and Security ACLs

Policy-Map Conditions	Security-ACL Conditions	Action
When the packet is in profile.	Permit specified packets.	Traffic is forwarded.
When the packet is out of profile and the out-of-profile action is to mark down the DSCP value.	Drop specified packets.	Traffic is dropped.
When the packet is out of profile and the out-of-profile action is to drop the packet.	Permit specified packets.	Traffic is dropped.
	Drop specified packets.	Traffic is dropped.

Configuring Classification Using Port Trust States


This section describes how to classify incoming traffic by using port trust states:

- Configuring the Trust State on Ports within the QoS Domain, on page 423
- Configuring the CoS Value for an Interface, on page 425
- Configuring Trusted Boundary, on page 426
- Enabling Pass-Through Mode, on page 428

Configuring the Trust State on Ports within the QoS Domain

Packets entering a QoS domain are classified at the edge of the QoS domain. When the packets are classified at the edge, the switch port within the QoS domain can be configured to one of the trusted states because there is no need to classify the packets at every switch within the QoS domain. Figure 68 shows a sample network topology.

Figure 68. Port Trusted States within the QoS Domain



Beginning in privileged EXEC mode, follow these steps to configure the port to trust the classification of the traffic that it receives:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface interface-id	Enter interface configuration mode, and specify the interface to be trusted. Valid interfaces include physical interfaces.

Step	Command	Purpose
3.	mls qos trust [cos dscp]	<p>Configure the port trust state.</p> <p>By default, the port is not trusted.</p> <p>The keywords have these meanings:</p> <p>cos—Classifies ingress packets with the packet CoS values. For tagged IP packets, the DSCP value of the packet is modified based on the CoS-to-DSCP map. The egress queue assigned to the packet is based on the packet CoS value.</p> <p>dscp—Classifies ingress packets with packet DSCP values. For non-IP packets, the packet CoS value is set to 0 for tagged packets; the default port CoS is used for untagged packets. Internally, the switch modifies the CoS value by using the DSCP-to-CoS map.</p> <p>Use the cos keyword if your network is composed of Ethernet LANs.</p> <p>Use the dscp keyword if your network is not composed of only Ethernet LANs and if you are familiar with sophisticated QoS features and implementations.</p> <p>For more information about this command, refer to the command reference for this release.</p>
4.	end	Return to privileged EXEC mode.
5.	show mls qos interface [interface-id] [policers]	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return a port to its untrusted state, use the **no mls qos trust** interface configuration command.

For information on how to change the default CoS value, see the “Configuring the CoS Value for an Interface” section on page 425. For information on how to configure the CoS-to-DSCP map, see the “Configuring the CoS-to-DSCP Map” section on page 436.

Configuring the CoS Value for an Interface

QoS assigns the CoS value specified with the **mls qos cos** interface configuration command to untagged frames received on trusted and untrusted ports.

Beginning in privileged EXEC mode, follow these steps to define the default CoS value of a port or to assign the default CoS to all incoming packets on the port:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface interface-id	<p>Enter interface configuration mode, and specify the interface to be trusted.</p> <p>Valid interfaces include physical interfaces.</p>

Step	Command	Purpose
3.	mls qos cos { <i>default-cos</i> override }	<p>Configure the default CoS value for the port.</p> <ul style="list-style-type: none"> For <i>default-cos</i>, specify a default CoS value to be assigned to a port. If the port is CoS trusted and packets are untagged, the default CoS value becomes the CoS value for the packet. The CoS range is 0 to 7. The default is 0. Use the override keyword to override the previously configured trust state of the incoming packets and to apply the default port CoS value to all incoming packets. By default, CoS override is disabled. <p>Use the override keyword when all incoming packets on certain ports deserve higher priority than packets entering from other ports. Even if a port was previously set to trust DSCP, this command overrides the previously configured trust state, and all the incoming CoS values are assigned the default CoS value configured with this command. If an incoming packet is tagged, the CoS value of the packet is modified with the default CoS of the port at the egress port.</p>
4.	end	Return to privileged EXEC mode.
5.	show mls qos interface	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no mls qos cos** {*default-cos* | **override**} interface configuration command.

Configuring Trusted Boundary

In a typical network, you connect a Cisco IP Phone to a switch port as shown in Figure 68. on page 424. Traffic sent from the telephone to the switch is typically marked with a tag that uses the 802.1Q header. The header contains the VLAN information and the CoS 3-bit field, which determines the priority of the packet. For most Cisco IP Phone configurations, the traffic sent from the telephone to the switch is trusted to ensure that voice traffic is properly prioritized over other types of traffic in the network. By using the **mls qos trust cos** interface configuration command, you can configure the switch port to which the telephone is connected to trust the CoS labels of all traffic received on that port.

In some situations, you also might connect a PC or workstation to the IP phone. In these cases, you can use the **switchport priority extend cos** interface configuration command to configure the telephone through the switch CLI to override the priority of the traffic received from the PC. With this command, you can prevent a PC from taking advantage of a high-priority data queue.

However, if a user bypasses the telephone and connects the PC directly to the switch, the CoS labels generated by the PC are trusted by the switch (because of the trusted CoS setting) and can allow misuse of high-priority queues. The trusted boundary feature solves this problem by using the CDP to detect the presence of a Cisco IP Phone (such as the Cisco IP Phone 7910, 7935, 7940, and 7960) on a switch port. If the telephone is not detected, the trusted boundary feature disables the trusted setting on the switch port and prevents misuse of a high-priority queue.

Beginning in privileged EXEC mode, follow these steps to configure trusted boundary on a switch port:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	cdp enable	Enable CDP globally. By default, it is enabled.
3.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be trusted. Valid interfaces include physical interfaces.
4.	cdp enable	Enable CDP on the interface. By default, CDP is enabled.
5.	mls qos trust device cisco-phone	Configure the Cisco IP Phone as a trusted device on the interface. You cannot enable both trusted boundary and auto-QoS (auto qos voip interface configuration command) at the same time; they are mutually exclusive.
6.	mls qos trust cos	Configure the port trust state to trust the CoS value of the ingress packet. By default, the port is not trusted. For more information on this command, refer to the command reference for this release.
7.	end	Return to privileged EXEC mode.
8.	show mls qos interface [<i>interface-id</i>] [policers]	Verify your entries.
9.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

When you enter the **no mls qos trust** interface configuration command, trusted boundary is not disabled. If this command is entered and the port is connected to a Cisco IP Phone, the port does not trust the classification of traffic that it receives. To disable trusted boundary, use the **no mls qos trust device** interface configuration command

If you enter the **mls qos cos override** interface configuration command, the port does not trust the classification of the traffic that it receives, even when it is connected to a Cisco IP Phone.

You cannot enable trusted boundary if auto-QoS is already enabled and vice-versa. If auto-QoS is enabled and a Cisco IP Phone is absent on a port, the port does not trust the classification of traffic that it receives.

Table 87 lists the port configuration when an IP phone is present or absent.

Table 87. Port Configurations When Trusted Boundary is Enabled

Port Configuration	When a Cisco IP Phone is Present	When a Cisco IP Phone is Absent
The port trusts the CoS value of the incoming packet.	The packet CoS value is trusted.	The packet CoS value is assigned the default CoS value.
The port trusts the DSCP value of the incoming packet.	The packet DSCP value is trusted.	For tagged non-IP packets, the packet CoS value is set to 0. For untagged non-IP packets, the packet CoS value is assigned the default CoS value.
The port assigns the default CoS value to incoming packets.	The packet CoS value is assigned the default CoS value.	The packet CoS value is assigned the default CoS value.

Enabling Pass-Through Mode

The switch assigns a CoS value of 0 to all incoming packets without modifying the packets. The switch offers best-effort service to each packet regardless of the packet contents or size and sends it from a single egress queue.

Beginning in privileged EXEC mode, follow these steps to enable pass-through mode:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface on which pass-through mode is enabled. Valid interfaces include physical interfaces.
3.	mls qos trust cos pass-through dscp	Enable pass-through mode. The interface is configured to trust the CoS value of the incoming packets and to send them without modifying the DSCP value.
4.	end	Return to privileged EXEC mode.
5.	show mls qos interface <i>[interface-id]</i>	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable pass-through mode, use the **no mls qos trust pass-through dscp** interface configuration command.

If you enter the **mls qos cos override** and the **mls qos trust [cos | dscp]** interface commands when pass-through mode is enabled, pass-through mode is disabled.

If you enter the **mls qos trust cos pass-through dscp** interface configuration command when the **mls qos cos override** and the **mls qos trust [cos | dscp]** interface commands are already configured, pass-through mode is disabled.

Configuring a QoS Policy

Configuring a QoS policy typically requires classifying traffic into classes, configuring policies applied to those traffic classes, and attaching policies to interfaces.

For background information, see the “Classification” section on page 413 and the “Policing and Marking” section on page 415.

This section contains this configuration information:

- Classifying Traffic by Using ACLs, on page 428
- Classifying Traffic by Using Class Maps, on page 432
- Classifying, Policing, and Marking Traffic by Using Policy Maps, on page 433

Classifying Traffic by Using ACLs

You can classify IP traffic by using IP standard or IP extended ACLs; you can classify Layer 2 traffic by using Layer 2 MAC ACLs.

Beginning in privileged EXEC mode, follow these steps to create an IP standard ACL for IP traffic:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	access-list <i>access-list-number</i> { permit remark } { <i>source source-wildcard</i> host <i>source</i> any }	<p>Create an IP standard ACL, repeating the command as many times as necessary.</p> <p>For <i>access-list-number</i>, enter the ACL number. The range is 1 to 99 and 1300 to 1999.</p> <p>Enter permit to specify whether to permit access if conditions are matched.</p> <p>Enter remark to specify an ACL entry comment up to 100 characters.</p> <p>Note: Deny statements are not supported for QoS ACLs. See the “Classification Based on QoS ACLs” section on page 414 for more details.</p> <p>The <i>source</i> is the source address of the network or host from which the packet is being sent, specified in one of three ways:</p> <ul style="list-style-type: none"> • The 32-bit quantity in dotted decimal format. • The keyword any as an abbreviation for <i>source</i> and <i>source-wildcard</i> of 0.0.0.0 255.255.255.255. You do not need to enter a source wildcard. • The keyword host as an abbreviation for <i>source</i> and <i>source-wildcard</i> of <i>source</i> 0.0.0.0. <p>(Optional) The <i>source-wildcard</i> variable applies wildcard bits to the source (see first bullet item).</p>
3.	end	Return to privileged EXEC mode.
4.	show access-lists	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

For more information about creating IP standard ACLs, see the “Guidelines for Applying ACLs to Physical Interfaces” section on page 375.

To delete an ACL, use the **no access-list** *access-list-number* global configuration command.

This example shows how to allow access for only those hosts on the two specified networks. The wildcard bits apply to the host portions of the network addresses. Any host with a source address that does not match the ACL statements is rejected.

```
Switch(config)# access-list 1 permit 192.5.255.0 0.0.0.255
Switch(config)# access-list 1 permit 36.0.0.0 0.0.0.255
```

Beginning in privileged EXEC mode, follow these steps to create an IP extended ACL for IP traffic:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	access-list <i>access-list-number</i> { permit remark } <i>protocol</i> { <i>source source-wildcard</i> host <i>source</i> any } [operator <i>port</i>] { <i>destination destination-wildcard</i> host <i>destination</i> any } [<i>operator</i> <i>port</i>] [dscp <i>dscp-value</i>] [time-range <i>time-range-name</i>]	<p>Create an IP extended ACL, repeating the command as many times as necessary.</p> <p>For <i>access-list-number</i>, enter the ACL number. The range is 100 to 199 and 2000 to 2699.</p> <p>Enter permit to permit access if conditions are matched.</p> <p>Enter remark to specify an ACL entry comment up to 100 characters.</p> <p>Note: Deny statements are not supported for QoS ACLs. See the “Classification Based on QoS ACLs” section on page 414 for more details.</p> <p>For <i>protocol</i>, enter the name or number of an IP protocol. Use the question mark (?) to see a list of available protocol keywords.</p> <p>For <i>source</i>, enter the network or host from which the packet is being sent. For <i>source-wildcard</i>, enter the wildcard bits by placing ones in the bit positions that you want to ignore. You specify the <i>source</i> and <i>source-wildcard</i> by using dotted decimal notation, by using the any keyword as an abbreviation for <i>source</i> 0.0.0.0 <i>source-wildcard</i> 255.255.255.255, or by using the host keyword for <i>source</i> 0.0.0.0.</p> <p>For <i>destination</i>, enter the network or host to which the packet is being sent. You have the same options for specifying the <i>destination</i> and <i>destination-wildcard</i> as those described for <i>source</i> and <i>source-wildcard</i>.</p> <p>Define a destination or source port.</p> <ul style="list-style-type: none"> • The <i>operator</i> can be only eq (equal). • If <i>operator</i> is after <i>source source-wildcard</i>, conditions match when the source port matches the defined port. • If <i>operator</i> is after <i>destination destination-wildcard</i>, conditions match when the destination port matches the defined port. • The <i>port</i> is a decimal number or name of a TCP or UDP port. The number can be from 0 to 65535. • Use TCP port names only for TCP traffic. • Use UDP port names only for UDP traffic. <p>Enter dscp to match packets with any of the 13 supported DSCP values (0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56) or use the question mark (?) to see a list of available values.</p> <p>The time-range keyword is optional. For information about this keyword, see the “Applying Time Ranges to ACLs” section on page 383.</p>
3.	end	Return to privileged EXEC mode.

Step	Command	Purpose
4.	show access-lists	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

For more information about creating IP extended ACLs, see the “Guidelines for Applying ACLs to Physical Interfaces” section on page 375.

To delete an ACL, use the **no access-list** *access-list-number* global configuration command.

This example shows how to create an ACL that permits only TCP traffic from the destination IP address 128.88.1.2 with TCP port number 25:

```
Switch(config)# access-list 102 permit tcp 0.0.0.0 255.255.255.255
128.88.1.2 0.0.0.0 eq 25
```

Beginning in privileged EXEC mode, follow these steps to create a Layer 2 MAC ACL for Layer 2 traffic:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	mac access-list extended <i>name</i>	Create a Layer 2 MAC ACL by specifying the name of the list. After entering this command, the mode changes to extended MAC ACL configuration.
3.	permit { any host <i>source MAC address</i> } { any host <i>destination MAC address</i> } [aarp amber appletalk dec-spanning decnet-iv diagnostic dsm etype-6000 etype-8042 lat lavc-sca mop-console mop-dump msdos mumps netbios vines-echo vines-ip xns-idp]	Enter permit to permit access if conditions are matched. Note: Deny statements are not supported for QoS ACLs. See the “Classification Based on QoS ACLs” section on page 414 for more details. For <i>source MAC address</i> , enter the MAC address of the host from which the packet is being sent. You specify this by using the any keyword to deny any source MAC address or by using the host keyword and the source in the hexadecimal format (H.H.H). For <i>destination MAC address</i> , enter the MAC address of the host to which the packet is being sent. You specify this by using the any keyword to deny any destination MAC address or by using the host keyword and the destination in the hexadecimal format (H.H.H). (Optional) You can also enter these options: aarp amber appletalk dec-spanning decnet-iv diagnostic dsm etype-6000 etype-8042 lat lavc-sca mop-console mop-dump msdos mumps netbios vines-echo vines-ip xns-idp (a non-IP protocol).
4.	end	Return to privileged EXEC mode.
5.	show access-lists [<i>number</i> <i>name</i>]	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

For more information about creating MAC extended ACLs, see the “Creating Named MAC Extended ACLs” section on page 386.

To delete an ACL, use the **no mac access-list extended *name*** global configuration command.

This example shows how to create a Layer 2 MAC ACL with a permit statement. The statement allows traffic from the host with MAC address 0001.0000.0001 to the host with MAC address 0002.0000.0001.

```
Switch(config)# mac access-list extended maclist1
Switch(config-ext-macl)# permit host 0001.0000.0001 host 0002.0000.0001
```

Classifying Traffic by Using Class Maps

You use the **class-map** global configuration command to isolate a specific traffic flow (or class) from all other traffic and to name it. The class map defines the criteria to use to match against a specific traffic flow to further classify it. Match statements can only include ACLs. The match criterion is defined with one match statement entered within the class-map configuration mode.

Note: You can also create class maps during policy map creation by using the **class** policy-map configuration command. For more information, see the “Classifying, Policing, and Marking Traffic by Using Policy Maps” section on page 433.

Beginning in privileged EXEC mode, follow these steps to create a class map and to define the match criterion to classify traffic:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	access-list <i>access-list-number</i> permit { <i>source source-wildcard</i> host <i>source</i> any } or access-list <i>access-list-number</i> {permit remark} <i>protocol</i> { <i>source source-wildcard</i> host <i>source</i> any } [<i>operator port</i>] { <i>destination destination-wildcard</i> host <i>destination</i> any } [<i>operator</i> <i>port</i>] [dscp <i>dscp-value</i>] [time-range <i>time-range-name</i>] or mac access-list extended <i>name</i> permit { any host <i>source MAC</i> <i>address</i> } { any host <i>destination</i> <i>MAC address</i> } [aarp amber dec- spanning decnet-iv diagnostic dsm etype-6000 etype-8042 lat lavc-sca mop-console mop- dump msdos mumps netbios vines-echo vines-ip xns-idp]	Create an IP standard or extended ACL for IP traffic or a Layer 2 MAC ACL for non-IP traffic, repeating the command as many times as necessary. For more information, see the “Guidelines for Applying ACLs to Physical Interfaces” section on page 375 and the “Classifying Traffic by Using ACLs” section on page 428. For more information on the mac access-list extended <i>name</i> command, see the “Creating Named MAC Extended ACLs” section on page 386. Note: Deny statements are not supported for QoS ACLs. See the “Classification Based on QoS ACLs” section on page 414 for more details.

Step	Command	Purpose
3.	class-map <i>class-map-name</i>	Create a class map, and enter class-map configuration mode. By default, no class maps are defined. For <i>class-map-name</i> , specify the name of the class map.
4.	match { access-group <i>acl-index</i> / access-group name <i>acl-name</i> ip dscp <i>dscp-list</i> }	Define the match criterion to classify traffic. By default, no match criterion is supported. Only one match criterion per class map is supported, and only one ACL per class map is supported. For access-group <i>acl-index</i> or access-group name <i>acl-name</i> , specify the number or name of the ACL created in Step 3. For ip dscp <i>dscp-list</i> , enter a list of up to eight IP DSCP values for each match statement to match against incoming packets. Separate each value with a space. The supported DSCP values are 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56.
5.	end	Return to privileged EXEC mode.
6.	show class-map [<i>class-map-name</i>]	Verify your entries.
7.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete an existing class map, use the **no class-map** *class-map-name* global configuration command. To remove a match criterion, use the **no match** {**access-group** *acl-index* | **name** *acl-name* | **ip dscp**} class-map configuration command.

This example shows how to configure the class map called *class1*. The *class1* has one match criterion, which is an ACL called *103*.

```
Switch(config)# access-list 103 permit any any tcp eq 80
Switch(config)# class-map class1
Switch(config-cmap)# match access-group 103
Switch(config-cmap)# end
Switch#
```

Classifying, Policing, and Marking Traffic by Using Policy Maps

A policy map specifies which traffic class to act on. Actions can include setting a specific DSCP value in the traffic class and specifying the traffic bandwidth limitations for each matched traffic class (policer) and the action to take when the traffic is out of profile (marking or dropping).

A policy map also has these characteristics:

- A policy map can contain multiple class statements, each with different match criteria and policers.
- A separate policy-map class can exist for each type of traffic received through an interface.

You can attach only one policy map per interface in the input direction.

Beginning in privileged EXEC mode, follow these steps to create a policy map:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	<p>access-list <i>access-list-number</i> permit {<i>source source-wildcard</i> host <i>source</i> any}</p> <p>or</p> <p>access-list <i>access-list-number</i> {permit remark} <i>protocol</i> {<i>source source-wildcard</i> host <i>source</i> any} [<i>operator port</i>] {<i>destination destination-wildcard</i> host <i>destination</i> any} [<i>operator port</i>] [<i>dscp dscp-value</i>] [time-range <i>time-range-name</i>]</p> <p>or</p> <p>mac access-list extended <i>name</i></p> <p>permit {any host <i>source MAC address</i>} {any host <i>destination MAC address</i>} [aarp amber appletalk dec-spanning decnet-iv diagnostic dsm etype-6000 etype-8042 lat lavc-sca mop-console mop-dump msdos mumps netbios vines-echo vines-ip xns-idp]</p>	<p>Create an IP standard or extended ACL for IP traffic or a Layer 2 MAC ACL for non-IP traffic, repeating the command as many times as necessary.</p> <p>For more information, see the “Classifying Traffic by Using ACLs” section on page 428.</p> <p>Note: Deny statements are not supported for QoS ACLs. See the “Classification Based on QoS ACLs” section on page 414 for more details.</p> <p>For more information on the mac access-list extended <i>name</i> command, see the “Creating Named MAC Extended ACLs” section on page 386.</p>
3.	policy-map <i>policy-map-name</i>	<p>Create a policy map by entering the policy map name, and enter policy-map configuration mode.</p> <p>By default, no policy maps are defined.</p> <p>The default behavior of a policy map is to set the DSCP to 0 if the packet is an IP packet and to set the CoS to 0 if the packet is tagged. No policing is performed.</p>
4.	class <i>class-map-name</i> [access-group <i>name</i> <i>acl-index-or-name</i>]	<p>Define a traffic classification, and enter policy-map class configuration mode.</p> <p>By default, no policy map class maps are defined.</p> <p>If a traffic class has already been defined by using the class-map global configuration command, specify its name for <i>class-map-name</i> in this command.</p> <p>For access-group <i>name</i> <i>acl-index-or-name</i>, specify the number or name of the ACL created in Step 2.</p> <p>Note: In a policy map, the class named <i>class-default</i> is not supported. The switch does not filter traffic based on the policy map defined by the class class-default policy-map configuration command.</p>
5.	set { ip dscp <i>new-dscp</i> }	<p>Classify IP traffic by setting a new value in the packet.</p> <p>For ip dscp <i>new-dscp</i>, enter a new DSCP value to be assigned to the classified traffic. The supported DSCP values are 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56.</p>

Step	Command	Purpose
6.	police <i>rate-bps burst-byte</i> [exceed-action { drop dscp <i>dscp-value</i> }]	<p>Define a policer for the classified traffic.</p> <p>You can configure up to 60 policers on ingress Gigabit-capable Ethernet ports.</p> <p>For <i>rate-bps</i>, specify average traffic rate in bits per second (bps). The range is 8 Mbps to 1000 Mbps for the Gigabit-capable Ethernet ports.</p> <p>For <i>burst-byte</i>, specify the normal burst size in bytes. The values supported on the Gigabit-capable Ethernet ports are 4096, 8192, 16348, 32768, 65536, 131072, 262144, and 524288.</p> <p>(Optional) Specify the action to take when the rates are exceeded. Use the exceed-action drop keywords to drop the packet. Use the exceed-action dscp <i>dscp-value</i> keywords to mark down the DSCP value and send the packet.</p>
7.	exit	Return to policy-map configuration mode.
8.	exit	Return to global configuration mode.
9.	interface <i>interface-id</i>	<p>Enter interface configuration mode, and specify the interface to attach to the policy map.</p> <p>Valid interfaces include physical interfaces.</p>
10.	service-policy input <i>policy-map-name</i>	<p>Apply specified policy map to the input of a particular interface.</p> <p>Only one policy map per interface per direction is supported.</p>
11.	end	Return to privileged EXEC mode.
12.	show policy-map [<i>policy-map-name</i> class <i>class-name</i>]	Verify your entries.
13.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete an existing policy map, use the **no policy-map** *policy-map-name* global configuration command. To delete an existing class map, use the **no class** *class-map-name* policy-map configuration command. To remove an assigned DSCP value, use the **no set ip dscp** *new-dscp* policy-map configuration command. To remove an existing policer, use the **no police** *rate-bps burst-byte* [**exceed-action** {**drop** | **dscp** *dscp-value*}] policy-map configuration command. To remove the policy map and interface association, use the **no service-policy input** *policy-map-name* interface configuration command.

For details about configuring policy maps and security ACLs on the same interface, see Table 86 on page 423.

This example shows how to create a policy map and attach it to an ingress interface. In the configuration, the IP standard ACL permits traffic from network 10.1.0.0. For traffic matching this classification, the DSCP value in the incoming packet is trusted. If the matched traffic exceeds an average traffic rate of 5000000 bps and a normal burst size of 8192 bytes, its DSCP is marked down to a value of 10 and sent.

```
Switch(config)# access-list 1 permit 10.1.0.0 0.0.255.255
Switch(config)# class-map ipclass1
Switch(config-cmap)# match access-group 1
Switch(config-cmap)# exit
Switch(config)# policy-map flow1t
```

```

Switch(config-pmap)# class ipclass1
Switch(config-pmap-c)# police 5000000 8192 exceed-action dscp 10
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# switchport mode access
Switch(config-if)# service-policy input flow1t

```

This example shows how to create a Layer 2 MAC ACL with two permit statements and attach it to an ingress interface. The first permit statement allows traffic from the host with MAC address 0001.0000.0001 destined for the host with MAC address 0002.0000.0001.

```

Switch(config)# mac access-list extended maclist1
Switch(config-ext-mac)# permit host 0001.0000.0001 host 0002.0000.0001
Switch(config-ext-mac)# exit
Switch(config)# mac access-list extended maclist2
Switch(config-ext-mac)# permit host 0001.0000.0003 host 0002.0000.0003
Switch(config-ext-mac)# exit
Switch(config)# class-map macclass1
Switch(config-cmap)# match access-group name maclist1
Switch(config-cmap)# exit
Switch(config)# policy-map macpolicy1
Switch(config-pmap)# class macclass1
Switch(config-pmap-c)# set ip dscp 56
Switch(config-pmap-c)# exit
Switch(config-pmap)# class macclass2 maclist2
Switch(config-pmap-c)# set ip dscp 48
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# switchport mode trunk
Switch(config-if)# mls qos trust cos
Switch(config-if)# service-policy input macpolicy1

```

Configuring CoS Maps

This section describes how to configure the CoS maps:

- Configuring the CoS-to-DSCP Map, on page 436
- Configuring the DSCP-to-CoS Map, on page 437

All the maps are globally defined.

Configuring the CoS-to-DSCP Map

You use the CoS-to-DSCP map to map CoS values in incoming packets to a DSCP value that QoS uses internally to represent the priority of the traffic.

Table 88 shows the default CoS-to-DSCP map.

Table 88. Default CoS-to-DSCP Map

CoS value	0	1	2	3	4	5	6	7
DSCP value	0	8	16	24	32	40	48	56

If these values are not appropriate for your network, you need to modify them.

Beginning in privileged EXEC mode, follow these steps to modify the CoS-to-DSCP map:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	mls qos map cos-dscp <i>dscp1...dscp8</i>	Modify the CoS-to-DSCP map. For <i>dscp1...dscp8</i> , enter 8 DSCP values that correspond to CoS values 0 to 7. Separate each DSCP value with a space. The supported DSCP values are 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56.
3.	end	Return to privileged EXEC mode.
4.	show mls qos maps cos-dscp	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default map, use the **no mls qos map cos-dscp** global configuration command.

This example shows how to modify and display the CoS-to-DSCP map:

```
Switch# configure terminal
Switch(config)# mls qos map cos-dscp 8 8 8 8 24 32 56 56
Switch(config)# end
Switch# show mls qos maps cos-dscp
```

```
Cos-dscp map:
  cos:  0  1  2  3  4  5  6  7
-----
  dscp:  8  8  8  8 24 32 56 56
```

Configuring the DSCP-to-CoS Map

You use the DSCP-to-CoS map to map DSCP values in incoming packets to a CoS value, which is used to select one of the four egress queues.

The switch supports these DSCP values: 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56.

Table 89 shows the default DSCP-to-CoS map.

Table 89. Default DSCP-to-CoS Map

DSCP values	0	8, 10	16, 18	24, 26	32, 34	40, 46	48	56
CoS values	0	1	2	3	4	5	6	7

If these values are not appropriate for your network, you need to modify them.

Beginning in privileged EXEC mode, follow these steps to modify the DSCP-to-CoS map:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	mls qos map dscp-cos <i>dscp-list</i> to <i>cos</i>	Modify the DSCP-to-CoS map. For <i>dscp-list</i> , enter up to 13 DSCP values separated by spaces. Then enter the to keyword. For <i>cos</i> , enter the CoS value to which the DSCP values correspond. The supported DSCP values are 0, 8, 10, 16, 18, 24, 26, 32, 34, 40, 46, 48, and 56. The CoS range is 0 to 7.
3.	end	Return to privileged EXEC mode.
4.	show mls qos maps dscp-cos	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default map, use the **no mls qos map dscp-cos** global configuration command.

This example shows how the DSCP values 26 and 48 are mapped to CoS value 7. For the remaining DSCP values, the DSCP-to-CoS mapping is the default.

```
Switch(config)# mls qos map dscp-cos 26 48 to 7
Switch(config)# exit
```

```
Switch# show mls qos maps dscp-cos
```

```
Dscp-cos map:
  dscp:  0  8 10 16 18 24 26 32 34 40 46 48 56
-----
   cos:  0  1  1  2  2  3  7  4  4  5  5  7  7
```

Configuring the Egress Queues

This section describes how to configure the egress queues:

- Configuring CoS Priority Queues, on page 438
- Configuring WRR Priority, on page 439
- Enabling the Expedite Queue and Configuring WRR Priority, on page 439

For more information about the egress queues, see the “Egress CoS Queues” section on page 417.

Configuring CoS Priority Queues

Beginning in privileged EXEC mode, follow these steps to configure the CoS priority queues:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	wrr-queue cos-map <i>qid</i> <i>cos1..cosn</i>	Specify the queue ID of the CoS priority queue. (Ranges are 1 to 4 where 1 is the lowest CoS priority queue.) Specify the CoS values that are mapped to the queue id. Default values are as follows: CoS ValueCoS Priority Queues 0, 11 2, 32 4, 53 6, 74
3.	end	Return to privileged EXEC mode.
4.	show wrr-queue cos-map	Display the mapping of the CoS priority queues.

To disable the new CoS settings and return to default settings, use the **no wrr-queue cos-map** global configuration command.

Configuring WRR Priority

Beginning in privileged EXEC mode, follow these steps to configure the WRR priority:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	wrr-queue bandwidth <i>weight1...weight4</i>	Assign WRR weights to the four CoS queues. These are the ranges for the WRR values: <ul style="list-style-type: none"> For <i>weight1</i>, <i>weight2</i>, and <i>weight3</i>, the range is 1 to 255. For <i>weight4</i>, the range is 0 to 255. When <i>weight4</i> is set to 0, queue 4 is configured as the expedite queue.
3.	end	Return to privileged EXEC mode.
4.	show wrr-queue bandwidth	Display the WRR bandwidth allocation for the CoS priority queues.

To disable the WRR scheduling and enable the strict priority scheduling, use the **no wrr-queue bandwidth** global configuration command.

To enable one of the queues as the expedite queue and to enable the WRR scheduling for the remaining queues, see the “Enabling the Expedite Queue and Configuring WRR Priority” section on page 439.

Enabling the Expedite Queue and Configuring WRR Priority

Beginning in privileged EXEC mode, follow these steps to enable the expedite queue (queue 4) and assign WRR priority to the remaining queues:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	wrr-queue bandwidth <i>weight1 weight2 weight3 0</i>	Configure queue 4 as the expedite queue and assign WRR weights to the remaining egress queues. The range of WRR weights for <i>weight1</i> , <i>weight2</i> , and <i>weight3</i> is 1 to 255.
3.	end	Return to privileged EXEC mode.
4.	show wrr-queue bandwidth	Display the WRR bandwidth allocation for the CoS priority queues.

Displaying Standard QoS Information

To display standard QoS information, use one or more of the privileged EXEC commands in Table 90:

Table 90. Commands for Displaying QoS Information

Command	Purpose
show class-map [<i>class-map-name</i>]	Display QoS class maps, which define the match criteria to classify traffic.
show policy-map [<i>policy-map-name</i> [class <i>class-name</i>]] ¹	Display QoS policy maps, which define classification criteria for incoming traffic.
show mls qos maps [cos-dscp dscp-cos] ¹	Display QoS mapping information. Maps are used to generate an internal DSCP value, which represents the priority of the traffic.
show mls qos interface [<i>interface-id</i>] [policers] ¹	Display QoS information at the interface level, including the configuration of the egress queues and the CoS-to-egress-queue map, which interfaces have configured policers, and ingress and egress statistics (including the number of bytes dropped).
show mls masks [qos security] ¹	Display details regarding the masks ¹ used for QoS and security ACLs.
show wrr-queue cos-map	Display the mapping of the CoS priority queues.
show wrr-queue bandwidth	Display the WRR bandwidth allocation for the CoS priority queues.


1. Access control parameters are called masks in the switch CLI commands and output.

Standard QoS Configuration Examples

This section shows a QoS migration path to help you quickly implement QoS features based on your existing network and planned changes to your network, as shown in Figure 69. It contains this information:

- QoS Configuration for the Existing Wiring Closet, on page 441
- QoS Configuration for the Intelligent Wiring Closet, on page 442

Figure 69. QoS Configuration Example Network



88288

QoS Configuration for the Existing Wiring Closet

The existing wiring closet in Figure 69 consists of existing Catalyst 2900 XL and 3500 XL switches. These switches are running Cisco IOS Release 12.0(5)XP or later, which supports the QoS-based IEEE 802.1p CoS values. QoS classifies frames by assigning priority-indexed CoS values to them and gives preference to higher-priority traffic.

Recall that on the Catalyst 2900 and 3500 XL switches, you can classify untagged (native) Ethernet frames at the ingress ports by setting a default CoS priority (**switchport priority default *default-priority-id*** interface configuration command) for each port. For IEEE 802.1Q frames with tag information, the priority value from the header frame is used. On the Catalyst 3524-PWR XL and 3548 XL switches, you can override this priority with the default value by using the **switchport priority default override** interface configuration command. For Catalyst 2950 and Catalyst 2900 XL switches and other 3500 XL models that do not have the override feature, the Catalyst 3550-12T switch at the distribution layer can override the 802.1p CoS value by using the **mls qos cos override** interface configuration command.

For the Catalyst 2900 and 3500 XL switches, CoS configures each transmit port (the egress port) with a normal-priority transmit queue and a high-priority transmit queue, depending on the frame tag or the port information. Frames in the normal-priority

queue are forwarded only after frames in the high-priority queue are forwarded. Frames that have 802.1p CoS values of 0 to 3 are placed in the normal-priority transmit queue while frames with CoS values of 4 to 7 are placed in the expedite (high-priority) queue.

QoS Configuration for the Intelligent Wiring Closet

The intelligent wiring closet in Figure 69 is composed of Catalyst 2950 switches. One of the switches is connected to a video server, which has an IP address of 172.20.10.16.

The object of this example is to prioritize the video traffic over all other traffic. To do so, a DSCP of 46 is assigned to the video traffic. This traffic is stored in queue 4, which is serviced more frequently than the other queues.

Beginning in privileged EXEC mode, follow these steps to configure the switch to prioritize video packets over all other traffic:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	access-list 1 permit 172.20.10.16	Define an IP standard ACL, and permit traffic from the video server at 172.20.10.16.
3.	class-map videoclass	Create a class map called <i>videoclass</i> , and enter class-map configuration mode.
4.	match access-group 1	Define the match criterion by matching the traffic specified by ACL 1.
5.	exit	Return to global configuration mode.
6.	policy-map videopolicy	Create a policy map called <i>videopolicy</i> , and enter policy-map configuration mode.
7.	class videoclass	Specify the class on which to act, and enter policy-map class configuration mode.
8.	set ip dscp 46	For traffic matching ACL 1, set the DSCP of incoming packets to 46.
9.	police 5000000 8192 exceed-action drop	Define a policer for the classified video traffic to drop traffic that exceeds 5-Mbps average traffic rate with an 8192-byte burst size.
10.	exit	Return to policy-map configuration mode.
11.	exit	Return to global configuration mode.
12.	interface gigabitethernet0/17	Enter interface configuration mode, and specify the ingress interface.
13.	service-policy input videopolicy	Apply the policy to the ingress interface.
14.	exit	Return to global configuration mode.
15.	wrr-queue bandwidth 1 2 3 4	Assign a higher WRR weight to queue 4.
16.	wrr-queue cos-map 4 6 7	Configure the CoS-to-egress-queue map so that CoS values 6 and 7 select queue 4.
17.	end	Return to privileged EXEC mode.
18.	show class-map videoclass show policy-map videopolicy show mls qos maps [cos-dscp dscp-cos]	Verify your entries.
19.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Chapter 25. Configuring EtherChannels

This chapter describes how to configure EtherChannel on the Layer 2 interfaces of a switch. EtherChannel provides fault-tolerant high-speed links between switches, routers, and servers. You can use it to increase the bandwidth between the wiring closets and the data center, and you can deploy it anywhere in the network where bottlenecks are likely to occur. EtherChannel provides automatic recovery for the loss of a link by redistributing the load across the remaining links. If a link fails, EtherChannel redirects traffic from the failed link to the remaining links in the channel without intervention.

Note: EtherChannel is supported on the external 10/100/1000 Mbps ports only. You can configure EtherChannel groups on the internal ports (ports 1-14 and ports 15 and 16). However, they will not exchange EtherChannel protocol messages and they will not function properly.

This chapter consists of these sections:


- Understanding EtherChannels, on page 443
- Configuring EtherChannels, on page 448
- Displaying EtherChannel, PAgP, and LACP Status, on page 455

Note: For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

Understanding EtherChannels

An EtherChannel consists of individual Fast Ethernet or Gigabit Ethernet links bundled into a single logical link as shown in Figure 70. The EtherChannel provides full-duplex bandwidth up to 800 Mbps (Fast EtherChannel) or 4 Gbps (Gigabit EtherChannel) between your switch and another switch or host.

Figure 70. Typical EtherChannel Configuration



Each EtherChannel can consist of up to eight compatibly configured Ethernet interfaces. All interfaces in each EtherChannel must be the same speed, and all must be configured as Layer 2 interfaces.

Note: The network device to which your switch is connected can impose its own limits on the number of interfaces in the EtherChannel. The number of EtherChannels is limited to six with eight ports per EtherChannel.

If a link within an EtherChannel fails, traffic previously carried over that failed link changes to the remaining links within the EtherChannel. A trap is sent for a failure, identifying the switch, the EtherChannel, and the failed link. Inbound broadcast and multicast packets on one link in an EtherChannel are blocked from returning on any other link of the EtherChannel.

Understanding Port-Channel Interfaces

When you create an EtherChannel for Layer 2 interfaces, a logical interface is dynamically created. You then manually assign an interface to the EtherChannel by using the **channel-group** interface configuration command.

Each EtherChannel has a logical port-channel interface numbered from 1 to 6.

When a port joins an EtherChannel, the physical interface for that port is shut down. When the port leaves the port-channel, its physical interface is brought up, and it has the same configuration as it had before joining the EtherChannel.

Note: Configuration changes made to the logical interface of an EtherChannel do not propagate to all the member ports of the channel.

Understanding the Port Aggregation Protocol and Link Aggregation Protocol

The Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP) facilitate the automatic creation of EtherChannels by exchanging packets between Ethernet interfaces. PAgP is a Cisco-proprietary protocol that can be run only on Cisco switches and on those switches licensed by licensed vendors to support PAgP. LACP is defined in IEEE 802.3ad and allows Cisco switches to manage Ethernet channels between switches that conform to the 802.3ad protocol.

By using one of these protocols, a switch learns the identity of partners capable of supporting either PAgP or LACP and learns the capabilities of each interface. It then dynamically groups similarly configured interfaces into a single logical link (channel or aggregate port); these interfaces are grouped based on hardware, administrative, and port parameter constraints. For example, PAgP groups the interfaces with the same speed, duplex mode, native VLAN, VLAN range, and trunking status and type. After grouping the links into an EtherChannel, PAgP adds the group to the spanning tree as a single switch port.

PAgP and LACP Modes

Table 91 shows the user-configurable EtherChannel modes for the **channel-group** interface configuration command. Switch interfaces exchange PAgP packets only with partner interfaces configured in the **auto** or **desirable** modes. Switch interfaces exchange LACP packets only with partner interfaces configured in the **active** or **passive** modes. Interfaces configured in the **on** mode do not exchange PAgP or LACP packets.

Table 91. EtherChannel Modes

Mode	Description
active	Places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending LACP packets.
auto	Places an interface into a passive negotiating state, in which the interface responds to PAgP packets it receives but does not start PAgP packet negotiation. This setting minimizes the transmission of PAgP packets.
desirable	Places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending PAgP packets.
on	Forces the interface into an EtherChannel without PAgP or LACP. With the on mode, a usable EtherChannel exists only when an interface group in the on mode is connected to another interface group in the on mode.
passive	Places an interface into a passive negotiating state, in which the interface responds to LACP packets that it receives, but does not start LACP packet negotiation. This setting minimizes the transmission of LACP packets.

Exchanging PAgP Packets

Both the **auto** and **desirable** PAgP modes allow interfaces to negotiate with partner interfaces to determine if they can form an EtherChannel based on criteria such as interface speed and, for Layer 2 EtherChannels, trunking state and VLAN numbers.

Interfaces can form an EtherChannel when they are in different PAgP modes as long as the modes are compatible. For example:

- An interface in the **desirable** mode can form an EtherChannel with another interface that is in the **desirable** or **auto** mode.
- An interface in the **auto** mode can form an EtherChannel with another interface in the **desirable** mode.

An interface in the **auto** mode cannot form an EtherChannel with another interface that is also in the **auto** mode because neither interface starts PAgP negotiation.

An interface in the **on** mode that is added to a port channel is forced to have the same characteristics as the already existing **on** mode interfaces in the channel.

If your switch is connected to a partner that is PAgP-capable, you can configure the switch interface for nonsilent operation by using the **non-silent** keyword. If you do not specify **non-silent** with the **auto** or **desirable** mode, silent mode is assumed.

The silent mode is used when the switch is connected to a device that is not PAgP-capable and seldom, if ever, sends packets. An example of a silent partner is a file server or a packet analyzer that is not generating traffic. In this case, running PAgP on a physical port connected to a silent partner prevents that switch port from ever becoming operational; however, the silent setting allows PAgP to operate, to attach the interface to a channel group, and to use the interface for transmission.

Note: An Etherchannel cannot be configured in both the PAgP and LACP modes.

Exchanging LACP Packets

Both the **active** and **passive LACP** modes allow interfaces to negotiate with partner interfaces to determine if they can form an EtherChannel based on criteria such as interface speed and, for Layer 2 EtherChannels, trunking state and VLAN numbers.

Interfaces can form an EtherChannel when they are in different LACP modes as long as the modes are compatible. For example:

- An interface in the **active** mode can form an EtherChannel with another interface that is in the **active** or **passive** mode.
- An interface in the **active** mode can form an EtherChannel with another interface in the **passive** mode.

An interface in the **passive** mode cannot form an EtherChannel with another interface that is also in the **passive** mode because neither interface starts LACP negotiation.

An interface in the **on** mode that is added to a port channel is forced to have the same characteristics as the already existing **on** mode interfaces in the channel.

Note: An Etherchannel cannot be configured in both the PAgP and LACP modes.

Caution: You should exercise care when setting the mode to on (manual configuration). All ports configured in the on mode are bundled in the same group and are forced to have similar characteristics. If the group is misconfigured, packet loss or spanning-tree loops might occur.

Physical Learners and Aggregate-Port Learners

Network devices are classified as PAgP physical learners or aggregate-port learners. A device is a physical learner if it learns addresses by physical ports and directs transmissions based on that knowledge. A device is an aggregate-port learner if it learns addresses by aggregate (logical) ports.

When a device and its partner are both aggregate-port learners, they learn the address on the logical port-channel. The device sends packets to the source by using any of the interfaces in the EtherChannel. With aggregate-port learning, it is not important on which physical port the packet arrives.

The switch uses source-MAC address distribution for a channel if it is connected to a physical learner even if you configure the switch for destination-MAC address distribution.

These frame distribution mechanisms are possible for frame transmission:

- Port selection based on the source-MAC address of the packet
- Port selection based on the destination- MAC address of the packet

The switch supports up to eight ports in a PAgP group.

PAgP and LACP Interaction with Other Features

The Dynamic Trunking Protocol (DTP) and Cisco Discovery Protocol (CDP) send and receive packets over the physical interfaces in the EtherChannel. Trunk ports send and receive PAgP and LACP protocol data units (PDUs) on the lowest numbered VLAN.

Spanning tree sends packets over a single physical interface in the EtherChannel. Spanning tree regards the EtherChannel as one port.

PAgP sends and receives PAgP PDUs only from interfaces that have PAgP enabled for the auto or desirable mode. LACP sends and receives LACP PDUs only from interfaces that have LACP enabled for the active or passive mode.

Understanding Load Balancing and Forwarding Methods

EtherChannel balances the traffic load across the links in a channel by randomly associating a newly-learned MAC address with one of the links in the channel.


With source-MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on the source-MAC address of the incoming packet. Therefore, to provide load balancing, packets from different hosts use different ports in the channel, but packets from the same host use the same port in the channel (and the MAC address learned by the switch does not change).

With destination-MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on the destination host's MAC address of the incoming packet. Therefore, packets to the same destination are forwarded over the same port, and packets to a different destination are sent on a different port in the channel. You configure the load balancing and forwarding method by using the **port-channel load-balance** global configuration command.

In Figure 71, multiple workstations are connected to a switch, and an EtherChannel connects the switch to the router. Source-based load balancing is used on the switch end of the EtherChannel to ensure that the switch efficiently uses the bandwidth of the router by distributing traffic from the workstation across the physical links. Since the router is a single MAC address device, it uses destination-based load balancing to efficiently spread the traffic to the workstations across the physical links in the EtherChannel.

Use the option that provides the greatest variety in your configuration. For example, if the traffic on a channel is going only to a single MAC address, using the destination-MAC address always chooses the same link in the channel; using source addresses or IP addresses might result in better load balancing.

Figure 71. Load Distribution and Forwarding Methods



74619

Configuring EtherChannels

These sections describe how to configure EtherChannel interfaces:

- Default EtherChannel Configuration, on page 449
- EtherChannel Configuration Guidelines, on page 449
- Configuring Layer 2 EtherChannels, on page 450
- Configuring EtherChannel Load Balancing, on page 452
- Configuring the PAgP Learn Method and Priority, on page 453
- Configuring the LACP Port Priority, on page 453
- Configuring Hot Standby Ports, on page 454
- Configuring the LACP System Priority, on page 454

Note: EtherChannel is supported on the external 10/100/1000 Mbps ports only.

Note: Make sure that the interfaces are correctly configured (see the “EtherChannel Configuration Guidelines” section on page 449).

Note: After you configure an EtherChannel, configuration changes applied to the port-channel interface apply to all the physical interfaces assigned to the port-channel

interface, and configuration changes applied to the physical interface affect only the interface where you apply the configuration.

Default EtherChannel Configuration

Table 92 shows the default EtherChannel configuration.

Table 92. Default EtherChannel Configuration

Feature	Default Setting
Channel groups	None assigned.
PAgP mode	No default.
PAgP learn method	Aggregate-port learning on all interfaces.
PAgP priority	128 on all interfaces. (Changing this value has no effect.)
LACP learn method	Aggregate-port learning on all interfaces.
LACP priority	32768 on all interfaces.
Load balancing	Load distribution on the switch is based on the source-MAC address of the incoming packet.

EtherChannel Configuration Guidelines

If improperly configured, some EtherChannel interfaces are automatically disabled to avoid network loops and other problems. Follow these guidelines to avoid configuration problems:

- Configure an EtherChannel with up to eight Ethernet interfaces of the same type.
- Configure all interfaces in an EtherChannel to operate at the same speeds and duplex modes.
- Configure EtherChannels only on external ports (17- 20).
- Enable all interfaces in an EtherChannel. An interface in an EtherChannel that is disabled by using the **shutdown** interface configuration command is treated as a link failure, and its traffic is transferred to one of the remaining interfaces in the EtherChannel.
- When a group is first created, all ports follow the parameters set for the first port to be added to the group. If you change the configuration of one of these parameters, you must also make the changes to all ports in the group:
 - Allowed-VLAN list
 - Spanning-tree path cost for each VLAN
 - Spanning-tree port priority for each VLAN
 - Spanning-tree Port Fast setting
- Do not configure a secure port as part of an EtherChannel.
- Before enabling 802.1X on the port, you must first remove it from the EtherChannel. If you try to enable 802.1X on an EtherChannel or on an active port in an EtherChannel, an error message appears, and 802.1X is not enabled. If you enable 802.1X on a not-yet active port of an EtherChannel, the port does not join the EtherChannel.
- An EtherChannel supports the same allowed range of VLANs on all the interfaces in a trunking Layer 2 EtherChannel. When configuring an interface for PAgP, if the allowed range of VLANs is not the same, the interfaces do not form an EtherChannel even when PAgP is set to the **auto** or **desirable** mode. When configuring an interface for LACP, if the allowed range of VLANs is not the same,

the interfaces do not form an EtherChannel even when LACP is set to the **active** or **passive** mode.

- Interfaces with different spanning-tree path costs can form an EtherChannel if they are otherwise compatibly configured. Setting different spanning-tree path costs does not, by itself, make interfaces incompatible for the formation of an EtherChannel.

Configuring Layer 2 EtherChannels

You configure Layer 2 EtherChannels by configuring the Ethernet interfaces with the **channel-group** interface configuration command, which creates the port-channel logical interface. You cannot put a Layer 2 interface into a manually created port-channel interface.

Note: Layer 2 interfaces must be connected and functioning for Cisco the software to create port-channel interfaces.

Beginning in privileged EXEC mode, follow these steps to assign a Layer 2 Ethernet interface to a Layer 2 EtherChannel:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify a physical interface to configure. Valid interfaces include physical interfaces. Up to eight interfaces of the same type and speed can be configured for the same group.

Step	Command	Purpose
3.	channel-group <i>channel-group-number</i> mode { auto [non-silent] desirable [non-silent] on } { active passive }	<p>Assign the interface to a channel group, and specify the PAgP or LACP mode.</p> <p>For <i>channel-group-number</i>, the range is 1 to 6. Each EtherChannel can have up to eight compatibly configured Ethernet interfaces.</p> <p>For mode, select one of these keywords:</p> <ul style="list-style-type: none"> • active—Enables LACP only if an LACP device is detected. It places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending LACP packets. • auto—Enables PAgP only if a PAgP device is detected. It places an interface into a passive negotiating state, in which the interface responds to PAgP packets it receives but does not start PAgP packet negotiation. • desirable—Unconditionally enables PAgP. It places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending PAgP packets. • on—Forces the interface to channel without PAgP. With the on mode, a usable EtherChannel exists only when an interface group in the on mode is connected to another interface group in the on mode. • non-silent—If your switch is connected to a partner that is PAgP-capable, you can configure the switch interface for nonsilent operation. You can configure an interface with the non-silent keyword for use with the auto or desirable mode. If you do not specify non-silent with the auto or desirable mode, silent is assumed. The silent setting is for connections to file servers or packet analyzers. This setting allows PAgP to operate, to attach the interface to a channel group, and to use the interface for transmission. • passive—Enables LACP on an interface and places it into a passive negotiating state, in which the interface responds to LACP packets that it receives, but does not start LACP packet negotiation. <p>For information on compatible PAgP and LACP modes for the switch and its partner, see the “PAgP and LACP Modes” section on page 444.</p>
4.	end	Return to privileged EXEC mode.
5.	show running-config	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove an interface from the EtherChannel group, use the **no channel-group** interface configuration command. If you delete the EtherChannel by using the **no interface port-channel** global configuration command without removing the physical

interfaces, the physical interfaces are shutdown. If you do not want the member physical interfaces to shut down, remove the physical interfaces before deleting the EtherChannel.

This example shows how to assign Gigabit Ethernet interfaces 0/17 and 0/20 with the PAgP mode **desirable**:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet0/17 -20
Switch(config-if)# channel-group 5 mode desirable
Switch(config-if)# end
```

Configuring EtherChannel Load Balancing

This section describes how to configure EtherChannel load balancing by using source-based or destination-based forwarding methods. For more information, see the “Understanding Load Balancing and Forwarding Methods” section on page 446.

Beginning in privileged EXEC mode, follow these steps to configure EtherChannel load balancing:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	port-channel load-balance {dst-mac src-mac}	<p>Configure an EtherChannel load-balancing method.</p> <p>The default is src-mac.</p> <p>Select one of these keywords to determine the load-distribution method:</p> <ul style="list-style-type: none"> • dst-mac—Load distribution is based on the destination-host MAC address of the incoming packet. Packets to the same destination are sent on the same port, but packets to different destinations are sent on different ports in the channel. • src-mac—Load distribution is based on the source-MAC address of the incoming packet. Packets from different hosts use different ports in the channel, but packets from the same host use the same port in the channel. <p>If the link partner to the switch is a physical learner, set the load-distribution method to one of these ways:</p> <ul style="list-style-type: none"> • If the channel-group interface configuration command is set to auto or desirable, the switch automatically uses the load distribution method based on the source-MAC address, regardless of the configured load-distribution method. • If the channel-group interface configuration command is set to on, set the load-distribution method based on the source-MAC address by using the port-channel load-balance src-mac global configuration command.
3.	end	Return to privileged EXEC mode.
4.	show etherchannel load-balance	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return EtherChannel load balancing to the default configuration, use the **no port-channel load-balance** global configuration command.

Configuring the PAgP Learn Method and Priority

Network devices are classified as PAgP physical learners or aggregate-port learners. A device is a physical learner if it learns addresses by physical ports and directs transmissions based on that knowledge. A device is an aggregate-port learner if it learns addresses by aggregate ports.

The switch supports address learning only on aggregate ports even though the **physical-port** keyword is provided in the command-line interface (CLI). The **pagp learn-method** and the **pagp port-priority** interface configuration command have no effect on the switch hardware.

Note: You should not set the learn method to **physical-port** because the switch is an aggregate-learning device.

If the link partner to the switch is a physical learner that has the **channel-group** interface configuration command set to **auto** or **desirable**, the switch automatically uses the load-distribution method based on the source MAC address, regardless of the configured load distribution method.

If the link partner to the switch is a physical learner that has the **channel-group** interface configuration command set to **on**, set the load-distribution method based on the source MAC address by using the **port-channel load-balance src-mac** global configuration command.

Configuring the LACP Port Priority

You can set the priority for each port in an EtherChannel that is configured for LACP by using the **lACP port-priority** privileged EXEC command. The range is from 1 to 65535. Beginning in privileged EXEC mode, follow these steps to configure the LACP port priority:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface for transmission.
3.	lACP port-priority <i>priority-value</i>	Select the LACP port priority value. For <i>priority-value</i> , the range is 1 to 65535. By default, the priority value is 32768. The lower the range, the more likely that the interface will be used for LACP transmission.
4.	end	Return to privileged EXEC mode.
5.	show running-config or show lACP <i>channel-group-number</i> internal	Verify your entries.
6.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Configuring Hot Standby Ports

When enabled, LACP tries to configure the maximum number of LACP-compatible ports in a channel, up to a maximum of 16 ports. Only eight LACP links can be active at one time. Any additional links are put in a hot standby state. If one of the active links becomes inactive, a link that is in hot standby mode becomes active in its place.

If more than eight links are configured for an EtherChannel group, the software determines which of the hot standby ports to make active based on:

- LACP port-priority
- Port ID

All ports default to the same port priority. You can change the port priority of LACP EtherChannel ports to specify which hot standby links become active first by using the **lACP port-priority** interface configuration command to set the port priority to a value lower than the default of 32768.

The hot standby ports that have lower port numbers become active in the channel first unless the port priority is configured to be a lower number than the default value of **32768**.

Note: If LACP is not able to aggregate all the ports that are compatible (for example, the remote system might have more restrictive hardware limitations), all the ports that cannot be actively included in the EtherChannel are put in hot standby state and are used only if one of the channeled ports fails.

Configuring the LACP System Priority

You can set the system priority for all of the EtherChannels that are configured for LACP by using the **lACP system-priority** privileged EXEC command. The range is from 1 to 65535.

Note: The **lACP system-priority** command is global. You cannot set a system priority for each LACP-configured channel separately.

We recommend using this command only when there are a combination of LACP-configured EtherChannels that are in both **active** and **standby** modes.

Beginning in privileged EXEC mode, follow these steps to configure the LACP system priority:

Step	Command	Purpose
1.	configure terminal	Enter global configuration mode.
2.	lacp system-priority <i>priority-value</i>	Select the LACP system priority value. For <i>priority-value</i> , the range is 1 to 65535. By default, the priority value is 32768. The lower the range, the higher the system priority. The switch with the lower system priority value determines which links between LACP partner switches are active and which are in standby for each LACP EtherChannel.
3.	end	Return to privileged EXEC mode.
4.	show running-config or show lacp <i>channel-group-number</i> internal	Verify your entries.
5.	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Displaying EtherChannel, PAgP, and LACP Status

You can use the privileged EXEC commands described in Table 93 to display EtherChannel, PAgP, and LACP status information:

Table 93. Commands for Displaying EtherChannel, PAgP, and LACP Status

Command	Description
show etherchannel [<i>channel-group-number</i>] { detail load-balance port port-channel summary }	Displays EtherChannel information in a detailed and one-line summary form. Also displays the load-balance or frame-distribution scheme, port, and port-channel information.
show pagp [<i>channel-group-number</i>] { counters internal neighbor } ¹	Displays PAgP information such as traffic information, the internal PAgP configuration, and neighbor information.
show lacp [<i>channel-group-number</i>] { counters internal neighbor } ²	Displays LACP information such as traffic information, the internal PAgP configuration, and neighbor information.

1. You can clear PAgP channel-group information and traffic filters by using the **clear pagp** [*channel-group-number*] [**counters**] | **counters**) privileged EXEC command.

2. You can clear LACP channel-group information and traffic filters by using the **clear lacp** [*channel-group-number*] [**counters**] | **counters**) privileged EXEC command.

For detailed information about the fields in the command outputs, refer to the command reference for this release.

Chapter 26. Troubleshooting

This chapter describes how to identify and resolve switch problems related to the Cisco IOS software. Depending on the nature of the problem, you can use the command-line interface (CLI) or the Cluster Management Suite (CMS) to identify and solve problems.

Note: For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release and the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference*.

This chapter consists of these sections:

- Preventing Autonegotiation Mismatches, on page 457
- Diagnosing Connectivity Problems, on page 457
- Using Debug Commands, on page 460
- Using the crashinfo File, on page 463
- Using the crashinfo File, on page 463

Preventing Autonegotiation Mismatches

The IEEE 802.3ab autonegotiation protocol manages the switch settings for speed (10 Mbps, 100 Mbps, and 1000 Mbps ports) and duplex (half or full). There are situations when this protocol can incorrectly align these settings, reducing performance. A mismatch occurs under these circumstances:

- A manually-set speed or duplex parameter is different from the manually set speed or duplex parameter on the connected port.
- A port is set to autonegotiate, and the connected port is set to full duplex with no autonegotiation.

To maximize switch performance and ensure a link, follow one of these guidelines when changing the settings for duplex and speed:

- Let both ports autonegotiate both speed and duplex.
- Manually set the speed and duplex parameters for the ports on both ends of the connection.

Note: If a remote device does not autonegotiate, configure the duplex settings on the two ports to match. The speed parameter can adjust itself even if the connected port does not autonegotiate.

Diagnosing Connectivity Problems

This section describes how to troubleshoot connectivity problems:

- Using Ping, on page 457
- Using Layer 2 Traceroute, on page 459

Using Ping

This section consists of this information:

- Understanding Ping, on page 458
- Executing Ping, on page 458

Understanding Ping

The switch supports IP ping, which you can use to test connectivity to remote hosts. Ping sends an echo request packet to an address and waits for a reply. Ping returns one of these responses:

- Normal response—The normal response (*hostname is alive*) occurs in 1 to 10 seconds, depending on network traffic.
- Destination does not respond—If the host does not respond, a *no-answer* message is returned.
- Unknown host—If the host does not exist, an *unknown host* message is returned.
- Destination unreachable—If the default gateway cannot reach the specified network, a *destination-unreachable* message is returned.
- Network or host unreachable—If there is no entry in the route table for the host or network, a *network or host unreachable* message is returned.

Executing Ping

If you attempt to ping a host in a different IP subnetwork, you must define a static route to the network.

Beginning in privileged EXEC mode, use this command to ping another device on the network from the switch:

Command	Purpose
<code>ping [ip] {host address}</code>	Ping a remote host through IP or by supplying the host name or network address.

Note: Though other protocol keywords are available with the **ping** command, they are not supported in this release.

This example shows how to ping an IP host:

```
Switch# ping 172.20.52.3
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echoes to 172.20.52.3, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

```
Switch#
```

Table 94 describes the possible ping character output.

Table 94. Ping Output Display Characters

Character	Description
!	Each exclamation point means receipt of a reply.
.	Each period means the network server timed out while waiting for a reply.
U	A destination unreachable error PDU was received.
C	A congestion experienced packet was received.
I	User interrupted test.
?	Unknown packet type.
&	Packet lifetime exceeded.

To terminate a ping session, enter the escape sequence (**Ctrl-^ X** by default). You enter the default by simultaneously pressing and releasing the **Ctrl**, **Shift**, and **6** keys, and then pressing the **X** key.

Using Layer 2 Traceroute

This section describes this information:

- Understanding Layer 2 Traceroute, on page 459
- Switches Supporting Layer 2 Traceroute, on page 459
- Usage Guidelines, on page 459
- Displaying the Physical Path, on page 460

Understanding Layer 2 Traceroute

The Layer 2 traceroute feature allows the switch to identify the physical path that a packet takes from a source device to a destination device. Layer 2 traceroute supports only unicast source and destination MAC addresses. It determines the path by using the MAC address tables of the switches in the path. When the switch detects a device in the path that does not support Layer 2 traceroute, the switch continues to send Layer 2 trace queries and lets them time out.

The switch can only identify the path from the source device to the destination device. It cannot identify the path that a packet takes from source host to the source device or from the destination device to the destination host.

Switches Supporting Layer 2 Traceroute

The Layer 2 traceroute feature is available on these switches:

- Catalyst 2940 switches
- Catalyst 2950 switches running Cisco IOS Release 12.1(12c)EA1 or later
- Catalyst 2955 switches
- Catalyst 3550 switches running Cisco IOS Release 12.1(12c)EA1 or later
- Catalyst 4000 switches running Catalyst software Release 6.2 or later for the supervisor engine
- Catalyst 5000 switches running Catalyst software Release 6.1 or later for the supervisor engine
- Catalyst 6000 switches running Catalyst software Release 6.1 or later for the supervisor engine
- Cisco Systems Intelligent Gigabit Ethernet Switch Module running 12.1(14)AY or later

Usage Guidelines

These are the Layer 2 traceroute usage guidelines:

- Cisco Discovery Protocol (CDP) must be enabled on all the devices in the network. For Layer 2 traceroute to functional properly, do not disable CDP.

For a list of switches that support Layer 2 traceroute, see the “Switches Supporting Layer 2 Traceroute” section on page 459. If any devices in the physical path are transparent to CDP, the switch cannot identify the path through these devices.

Note: For more information about enabling CDP, see Chapter 18 “Configuring CDP.”

- A switch is reachable from another switch when you can test connectivity by using the **ping** privileged EXEC command. All switches in the physical path must be reachable from each other.
- The maximum number of hops identified in the path is ten.
- You can enter the **traceroute mac** or the **traceroute mac ip** privileged EXEC command on a switch that is not in the physical path from the source device to the destination device. All switches in the path must be reachable from this switch.
- The **traceroute mac** command output shows the Layer 2 path only when the specified source and destination MAC addresses belong to the same VLAN. If you specify source and destination MAC addresses that belong to different VLANs, the Layer 2 path is not identified, and an error message appears.
- If you specify a multicast source or destination MAC address, the path is not identified, and an error message appears.
- If the source or destination MAC address belongs to multiple VLANs, you must specify the VLAN to which both the source and destination MAC addresses belong. If the VLAN is not specified, the path is not identified, and an error message appears.
- The **traceroute mac ip** command output shows the Layer 2 path when the specified source and destination IP addresses belong to the same subnet. When you specify the IP addresses, the switch uses Address Resolution Protocol (ARP) to associate the IP addresses with the corresponding MAC addresses and the VLAN IDs.
 - If an ARP entry exists for the specified IP address, the switch uses the associated MAC address and identifies the physical path.
 - If an ARP entry does not exist, the switch sends an ARP query and tries to resolve the IP address. If the IP address is not resolved, the path is not identified, and an error message appears.
- When multiple devices are attached to one port through hubs (for example, multiple CDP neighbors are detected on a port), the Layer 2 traceroute feature is not supported. When more than one CDP neighbor is detected on a port, the Layer 2 path is not identified, and an error message appears.
- This feature is not supported in Token Ring VLANs.

Displaying the Physical Path

You can display physical path that a packet takes from a source device to a destination device by using one of these privileged EXEC commands:

- **traceroute mac** [**interface** *interface-id*] {*source-mac-address*} [**interface** *interface-id*] {*destination-mac-address*} [**vlan** *vlan-id*] [**detail**]
- **traceroute mac ip** {*source-ip-address* / *source-hostname*}{*destination-ip-address* / *destination-hostname*} [**detail**]

For more information, refer to the command reference for this release.

Using Debug Commands

This section explains how you use the **debug** commands to diagnose and resolve internetworking problems. It contains this information:

- Enabling Debugging on a Specific Feature, on page 461
- Enabling All-System Diagnostics, on page 461
- Redirecting Debug and Error Message Output, on page 462

- Using the debug autoqos Command, on page 462

Caution: Because debugging output is assigned high priority in the CPU process, it can render the system unusable. For this reason, use debug commands only to troubleshoot specific problems or during troubleshooting sessions with technical support representatives. It is best to use debug commands during periods of lower network traffic and fewer users. Debugging during these periods decreases the likelihood that increased debug command processing overhead will affect system use.

Note: For complete syntax and usage information for specific **debug** commands, refer to the command reference for this release.

Enabling Debugging on a Specific Feature

All **debug** commands are entered in privileged EXEC mode, and most **debug** commands take no arguments. For example, beginning in privileged EXEC mode, enter this command to enable the debugging for EtherChannel:

```
Switch# debug etherchannel
```

The switch continues to generate output until you enter the **no** form of the command.

If you enable a **debug** command and no output appears, consider these possibilities:

- The switch might not be properly configured to generate the type of traffic that you want to monitor. Use the **show running-config** command to verify the configuration.
- Even if the switch is properly configured, it might not generate the type of traffic that you want to monitor during the particular period that debugging is enabled. Depending on the feature you are debugging, you can use commands such as the TCP/IP **ping** command to generate network traffic.

To disable debugging of SPAN, enter this command in privileged EXEC mode:

```
Switch# no debug etherchannel
```

Alternately, in privileged EXEC mode, you can enter the **undebug** form of the command:

```
Switch# undebug etherchannel
```

To display the state of each debugging option, enter this command in privileged EXEC mode:

```
Switch# show debugging
```

Enabling All-System Diagnostics

Beginning in privileged EXEC mode, enter this command to enable all-system diagnostics:

```
Switch# debug all
```

Caution: Because debugging output takes priority over other network traffic, and because the **debug all** privileged EXEC command generates more output than any other debug command, it can severely diminish switch performance or even render it unusable. In virtually all cases, it is best to use more specific debug commands.

The **no debug all** privileged EXEC command disables all diagnostic output. Using the **no debug all** command is a convenient way to ensure that you have not accidentally left any **debug** commands enabled.

Redirecting Debug and Error Message Output

By default, the network server sends the output from **debug** commands and system error messages to the service port. If you use this default, you can use a virtual terminal connection to monitor debug output instead of connecting to the service port.

Possible destinations include the service port, virtual terminals, internal buffer, and UNIX hosts running a syslog server. The syslog format is compatible with 4.3 Berkeley Standard Distribution (BSD) UNIX and its derivatives.

Note: Be aware that the debugging destination you use affects system overhead. Logging messages to the service port produces very high overhead, whereas logging messages to a virtual terminal produces less overhead. Logging messages to a syslog server produces even less, and logging to an internal buffer produces the least overhead of any method.

For more information about system message logging, see Chapter 21 “Configuring System Message Logging.”

Using the debug autoqos Command

You can use the **debug autoqos** privileged EXEC command to display quality of service (QoS) commands that are automatically generated when automatic-QoS (auto-QoS) is enabled.

Beginning in privileged EXEC mode, follow these steps to display the QoS commands and enable auto-QoS for voice over IP (VoIP) within a QoS domain:

Step	Command	Purpose
1.	debug autoqos	Enable debugging for auto-QoS. When debugging is enabled, the switch displays the QoS commands that are automatically generated when auto-QoS is enabled or disabled.
2.	configure terminal	Enter global configuration mode.
3.	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface that is connected to a Cisco IP Phone. You also can specify the uplink interface that is connected to another switch or router in the interior of the network.
4.	auto qos voip {cisco-phone trust}	Enable auto-QoS. The keywords have these meanings: <ul style="list-style-type: none"> cisco-phone—If the interface is connected to a Cisco IP Phone, the QoS labels of incoming packets are trusted only when the IP phone is detected. trust—The uplink interface is connected to a trusted switch or router, and the VoIP classification in the ingress packet is trusted.

Step	Command	Purpose
5.	end	Return to privileged EXEC mode.
6.	show auto qos interface <i>interface-id</i>	Verify your entries. This command displays the auto-QoS configuration that was initially applied; it does not display any user changes to the configuration that might be in effect.

For more information about auto-QoS, see the “Configuring Auto-QoS” section on page 417.

This example shows how to display the QoS commands that are automatically generated when auto-QoS is enabled:

```
Switch# debug autoqos
AutoQoS debugging is on
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet0/17
Switch(config-if)# auto qos voip cisco-phone
```

Using the crashinfo File

The crashinfo file saves information that helps technical support representatives to debug problems that caused the software image to fail (crash). The switch writes the crash information to the service port at the time of the failure, and the file is created the next time you boot the image after the failure (instead of while the system is failing).

The information in the file includes the software image name and version that failed, a dump of the processor registers, and a stack trace. You can give this information to the technical support representative by using the **show tech-support** privileged EXEC command.

All crashinfo files are kept in this directory on the flash file system:

flash:/crashinfo/crashinfo_*n* where *n* is a sequence number.

Each new crashinfo file that is created uses a sequence number that is larger than any previously existing sequence number, so the file with the largest sequence number describes the most recent failure. Version numbers are used instead of a timestamp because the switches do not include a real-time clock. You cannot change the name of the file that the system will use when it creates the file. However, after the file is created, you can use the **rename** privileged EXEC command to rename it, but the contents of the renamed file will not be displayed by the **show stacks** or the **show tech-support** privileged EXEC command. You can delete crashinfo files by using the **delete** privileged EXEC command.

You can display the most recent crashinfo file (that is, the file with the highest sequence number at the end of its filename) by entering the **show stacks** or the **show tech-support** privileged EXEC command. You also can access the file by using any command that can copy or display files, such as the **more** or the **copy** privileged EXEC command.

Appendix A. Supported MIBs

This appendix lists the supported management information base (MIBs) for this release. It contains these sections:

- MIB List, on page 465
- Accessing the MIB Files, on page 466

MIB List

- BRIDGE-MIB (RFC1493)
- CISCO-2900-MIB
- CISCO-BULK-FILE-MIB
- CISCO-CDP-MIB
- CISCO-CLUSTER-MIB
- CISCO-CONFIG-MAN-MIB
- CISCO-ENTITY-VENDORTYPE-OID-MIB
- CISCO-ENVMON-MIB
- CISCO-FLASH-MIB
- CISCO-FTP-CLIENT-MIB
- CISCO-IGMP-FILTER-MIB
- CISCO-IMAGE-MIB
- CISCO-IP-STAT-MIB
- CISCO-MAC-NOTIFICATION-MIB
- CISCO-MEMORY-POOL-MIB
- CISCO-PAGP-MIB
- CISCO-PING-MIB
- CISCO-PORT-SECURITY-MIB
- CISCO-PROCESS-MIB
- CISCO-PRODUCTS-MIB
- CISCO-RTTMON-MIB (subsystems supported: sub_rtt_rmon and sub_rtt_rmonlib)
- CISCO-SMI
- CISCO-STACKMAKER-MIB
- CISCO-STP-EXTENSIONS-MIB
- CISCO-SYSLOG-MIB
- CISCO-TC
- CISCO-TCP-MIB
- CISCO-VLAN-MEMBERSHIP-MIB
- CISCO-VTP-MIB
- ENTITY-MIB
- IANAifType-MIB
- IF-MIB (RFC 1573)

- OLD-CISCO-CHASSIS-MIB
- OLD-CISCO-CPU-MIB
- OLD-CISCO-INTERFACES-MIB
- OLD-CISCO-IP-MIB
- OLD-CISCO-MEMORY-MIB
- OLD-CISCO-SYSTEM-MIB
- OLD-CISCO-TCP-MIB
- OLD-CISCO-TS-MIB
- RFC1213-MIB
- RFC1398-MIB
- RMON-MIB (RFC 1757)
- RS-232-MIB
- SNMPv2-MIB
- SNMPv2-SMI
- SNMPv2-TC
- TCP-MIB
- UDP-MIB

Accessing the MIB Files

To access the Management Information Base (MIBs) for the CIGESM go to the IBM web site: www.ibm.com/support. Click on the **Download and Driver** icon. This will bring up a search web page. In the **Search** box type in Cisco and then click **submit**. A web page with a description of the latest software drivers and other pertinent information for the CIGESM is displayed. Click on this text and you will be directed to a web page listing software and pertinent information available for the CIGESM. Find the MIBs and then right click on this. This brings up a drop-down menu next to the file. Select **Save target as...** to save the zipped file to your hard disk. You will then need PKUNZIP to expand the files on your computer.

Appendix B. Working with the Cisco IOS File System, Configuration Files, and Software Images

This appendix describes how to manipulate the flash file system, how to copy configuration files, and how to archive (upload and download) software images on the switch.

Note: For complete syntax and usage information for the commands used in this chapter, refer to the switch command reference for this release and the *Cisco Systems Intelligent Gigabit Ethernet Switch Module for the IBM eServer BladeCenter System Command Reference*.

This appendix consists of these sections:

- Working with the Flash File System, on page 467
- Working with Configuration Files, on page 474
- Working with Software Images, on page 484

Working with the Flash File System

The flash file system on your switch provides several commands to help you manage software image and configuration files.

The flash file system is a single flash device on which you can store files. This flash device is called *flash*.

This section contains this information:

- Displaying Available File Systems, on page 467
- Setting the Default File System, on page 468
- Displaying Information about Files on a File System, on page 469
- Changing Directories and Displaying the Working Directory, on page 469
- Creating and Removing Directories, on page 469
- Copying Files, on page 470
- Deleting Files, on page 471
- Creating, Displaying, and Extracting tar Files, on page 471
- Displaying the Contents of a File, on page 473

Displaying Available File Systems

To display the available file systems on your switch, use the **show file systems** privileged EXEC command as shown in this example:

```
Switch# show file systems
File Systems:
```

	Size(b)	Free(b)	Type	Flags	Prefixes
*	16128000	11118592	flash	rw	flash:
	16128000	11118592	unknown	rw	zflash:
	32768	26363	nvr	rw	nvr:
	-	-	network	rw	tftp:
	-	-	opaque	rw	null:
	-	-	opaque	rw	system:

-	-	opaque	ro	xmodem:
-	-	opaque	ro	ymodem:
-	-	network	rw	rcp:
-	-	network	rw	ftp:

Table 95. *show file systems* Field Descriptions

Field	Value
Size(b)	Amount of memory in the file system in bytes.
Free(b)	Amount of free memory in the file system in bytes.
Type	Type of file system. flash —The file system is for a flash memory device. nvr am —The file system is for a nonvolatile RAM (NVRAM) device. opaque —The file system is a locally generated <i>pseudo</i> file system (for example, the <i>system</i>) or a download interface, such as brimux. unknown —The file system is an unknown type.
Flags	Permission for file system. ro —read-only. rw —read/write. wo —write-only.
Prefixes	Alias for file system. bs :—Read-only file system; stores the boot loader image. vb :—Stores the boot environment variables. flash :—Flash file system. nvr am :—NVRAM. null :—Null destination for copies. You can copy a remote file to null to determine its size. rcp :—Remote Copy Protocol (RCP) network server. system :—Contains the system memory, including the running configuration. tftp :—Trivial File Transfer Protocol (TFTP) network server. xmodem :—Obtain the file from a network machine by using the XMODEM protocol. ymodem :—Obtain the file from a network machine by using the YMODEM protocol. zflash :—Read-only file decompression file system, which mirrors the contents of the flash file system.

Setting the Default File System

You can specify the file system or directory that the system uses as the default file system by using the **cd *filesystem*:** privileged EXEC command. You can set the default file system to omit the *filesystem*: argument from related commands. For example, for all privileged EXEC commands that have the optional *filesystem*: argument, the system uses the file system specified by the **cd** command.

By default, the default file system is *flash*:

You can display the current default file system as specified by the **cd** command by using the **pwd** privileged EXEC command.

Displaying Information about Files on a File System

You can view a list of the contents of a file system before manipulating its contents. For example, before copying a new configuration file to flash memory, you might want to verify that the file system does not already contain a configuration file with the same name. Similarly, before copying a flash configuration file to another location, you might want to verify its filename for use in another command.

To display information about files on a file system, use one of the privileged EXEC commands in Table 96:

Table 96. Commands for Displaying Information About Files

Command	Description
dir [<i>all</i>] [<i>filesystem:</i>][<i>filename</i>]	Display a list of files on a file system.
show file systems	Display more information about each of the files on a file system.
show file information <i>file-url</i>	Display information about a specific file.
show file descriptors	Display a list of open file descriptors. File descriptors are the internal representations of open files. You can use this command to see if another user has a file open.

Changing Directories and Displaying the Working Directory

Beginning in privileged EXEC mode, follow these steps to change directories and display the working directory.

Step	Command	Purpose
1.	dir <i>filesystem:</i>	Display the directories on the specified file system. For <i>filesystem:</i> , use flash: for the system board flash device.
2.	cd <i>new_configs</i>	Change to the directory of interest. The command example shows how to change to the directory named <i>new_configs</i> .
3.	pwd	Display the working directory.

Creating and Removing Directories

Beginning in privileged EXEC mode, follow these steps to create and remove a directory:

Step	Command	Purpose
1.	dir <i>filesystem:</i>	Display the directories on the specified file system. For <i>filesystem:</i> , use flash: for the system board flash device.
2.	mkdir <i>old_configs</i>	Create a new directory. The command example shows how to create the directory named <i>old_configs</i> . Directory names are case sensitive. Directory names are limited to 45 characters between the slashes (/); the name cannot contain control characters, spaces, deletes, slashes, quotes, semicolons, or colons.
3.	dir <i>filesystem:</i>	Verify your entry.

To delete a directory with all its files and subdirectories, use the **delete /force /recursive** *filesystem:/file-url* privileged EXEC command.

Use the **/recursive** keyword to delete the named directory and all subdirectories and the files contained in it. Use the **/force** keyword to suppress the prompting that confirms a deletion of each file in the directory. You are prompted only once at the beginning of this deletion process. Use the **/force** and **/recursive** keywords for deleting old software images that were installed by using the **archive download-sw** command but are no longer needed.

For *filesystem*, use **flash:** for the system board flash device. For *file-url*, enter the name of the directory to be deleted. All the files in the directory and the directory are removed.

Caution: When files and directories are deleted, their contents cannot be recovered.

Copying Files

To copy a file from a source to a destination, use the **copy [/erase] source-url destination-url** privileged EXEC command. For the source and destination URLs, you can use **running-config** and **startup-config** keyword shortcuts. For example, the **copy running-config startup-config** command saves the currently running configuration file to the NVRAM section of flash memory to be used as the configuration during system initialization.

You can also copy to and from special file systems (**xmodem:**, **ymodem:**) as the source or destination for the file from a network machine that uses the XMODEM or YMODEM protocol.

Network file system URLs include **ftp:**, **rnp:**, and **tftp:** and have these syntaxes:

File Transfer Protocol (FTP)—**ftp:***[/username [:password]@location]/directory/filename*

Remote Copy Protocol (RCP)—**rnp:***[/username@location]/directory/filename*

Trivial File Transfer Protocol (TFTP)—**tftp:***[/location]/directory/filename*

Local writable file systems include flash:.

Some invalid combinations of source and destination exist. Specifically, you cannot copy these combinations:

- From a running configuration to a running configuration
- From a startup configuration to a startup configuration
- From a device to the same device (for example, the **copy flash: flash:** command is invalid)

For specific examples of using the **copy** command with configuration files, see the “Working with Configuration Files” section on page 474.

To copy software images either by downloading a new version or uploading the existing one, use the **archive download-sw** or the **archive upload-sw** privileged EXEC command. For more information, see the “Working with Software Images” section on page 484.

Deleting Files

When you no longer need a file on a flash memory device, you can permanently delete it. To delete a file or directory from a specified flash device, use the **delete** [**/force**] [**/recursive**] [*filesystem:*]*/file-url* privileged EXEC command.

Use the **/recursive** keyword for deleting a directory and all subdirectories and the files contained in it. Use the **/force** keyword to suppress the prompting that confirms a deletion of each file in the directory. You are prompted only once at the beginning of this deletion process. Use the **/force** and **/recursive** keywords for deleting old software images that were installed by using the **archive download-sw** command but are no longer needed.

If you omit the *filesystem:* option, the switch uses the default device specified by the **cd** command. For *file-url*, you specify the path (directory) and the name of the file to be deleted.

If you attempt to delete the file specified by the CONFIG_FILE or BOOT environment variable, the system prompts you to confirm the deletion. If you attempt to delete the last valid system image specified in the BOOT environment variable, the system prompts you to confirm the deletion.

Caution: When files are deleted, their contents cannot be recovered.

This example shows how to delete the file *myconfig* from the default flash memory device:

```
Switch# delete myconfig
```

Creating, Displaying, and Extracting tar Files

You can create a tar file and write files into it, list the files in a tar file, and extract the files from a tar file as described in the next sections.

Creating a tar File

To create a tar file and write files into it, use the privileged EXEC command:

```
archive tar /create destination-url flash:/file-url
```

For *destination-url*, specify the destination URL alias for the local or network file system and the name of the tar file to create. These options are supported:

- For the local flash file system, the syntax is **flash:**

- For the File Transfer Protocol (FTP), the syntax is
ftp:[*//username[:password]@location*]/*directory*]/*tar-filename.tar*
- For the Remote Copy Protocol (RCP), the syntax is
rcp:[*//username@location*]/*directory*]/*tar-filename.tar*
- For the Trivial File Transfer Protocol (TFTP), the syntax is
tftp:[*//location*]/*directory*]/*tar-filename.tar*

The *tar-filename.tar* is the tar file to be created.

For **flash:*/file-url***, specify the location on the local flash file system from which the new tar file is created. You can also specify an optional list of files or directories within the source directory to write to the new tar file. If none are specified, all files and directories at this level are written to the newly created tar file.

This example shows how to create a tar file. This command writes the contents of the *new-configs* directory on the local flash device to a file named *saved.tar* on the TFTP server at 172.20.10.30:

```
Switch# archive tar /create tftp:172.20.10.30/saved.tar flash:/new-configs
```

Displaying the Contents of a tar File

To display the contents of a tar file on the screen, use this privileged EXEC command:

```
archive tar /table source-url
```

For *source-url*, specify the source URL alias for the local or network file system. These options are supported:

- For the local flash file system, the syntax is
flash:
- For the File Transfer Protocol (FTP), the syntax is
ftp:[*//username[:password]@location*]/*directory*]/*tar-filename.tar*
- For the Remote Copy Protocol (RCP), the syntax is
rcp:[*//username@location*]/*directory*]/*tar-filename.tar*
- For the Trivial File Transfer Protocol (TFTP), the syntax is
tftp:[*//location*]/*directory*]/*tar-filename.tar*

The *tar-filename.tar* is the tar file to display.

You can also limit the display of the files by specifying an optional list of files or directories after the tar file; then only these files are displayed. If none are specified, all files and directories are displayed.

This example shows how to display the contents of the *cigesm-i6q4l2-mz.121-21.EA1.tar* file that is in flash memory:

```
Switch# archive tar /table flash:cigesm-i6q4l2-mz.121-21.EA1.tar
info (219 bytes)
cigesm-i6q4l2-mz.121-21.EA1/ (directory)
cigesm-i6q4l2-mz.121-21.EA1/html/ (directory)
cigesm-i6q4l2-mz.121-21.EA1/html/foo.html (0 bytes)
cigesm-i6q4l2-mz.121-21.EA1/cigesm-i6q4l2-mz.121-21.EA1.bin (610856 bytes)
cigesm-i6q4l2-mz.121-21.EA1/info (219 bytes)
info.ver (219 bytes)
```

This example shows how to display only the *cigesm-i6q4l2-mz.121-21.EA1/html* directory and its contents:

```
Switch# archive tar /table flash:cigesm-i6q412-mz.121-21.EA1/html
cigesm-i6q412-mz.121-21.EA1/html/ (directory)
cigesm-i6q412-mz.121-21.EA1/html/foo.html (0 bytes)
```

Extracting a tar File

To extract a tar file into a directory on the flash file system, use the privileged EXEC command:

```
archive tar /xtract source-url flash:/file-url
```

For *source-url*, specify the source URL alias for the local or network file system. These options are supported:

- For the local flash file system, the syntax is **flash:**
- For the File Transfer Protocol (FTP), the syntax is **ftp:[//username[:password]@location]/directory/tar-filename.tar**
- For the Remote Copy Protocol (RCP), the syntax is **rcp:[//username@location]/directory/tar-filename.tar**
- For the Trivial File Transfer Protocol (TFTP), the syntax is **tftp:[//location]/directory/tar-filename.tar**

The *tar-filename.tar* is the tar file from which to extract files.

For **flash:/file-url**, specify the location on the local flash file system into which the tar file is extracted. You can also specify an optional list of files or directories within the tar file for extraction. If none are specified, all files and directories are extracted.

This example shows how to extract the contents of a tar file located on the TFTP server at 172.20.10.30. This command extracts just the *new-configs* directory into the root directory on the local flash file system. The remaining files in the *saved.tar* file are ignored.

```
Switch# archive tar /xtract tftp://172.20.10.30/saved.tar flash:/new-configs
```

Displaying the Contents of a File

To display the contents of any readable file, including a file on a remote file system, use the **more [/ascii | /binary | /ebcdic] file-url** privileged EXEC command:

This example shows how to display the contents of a configuration file on a TFTP server:

```
Switch# more tftp://serverA/hampton/savedconfig
!
! Saved configuration on server
!
version 11.3
service timestamps log datetime localtime
service linenumber
service udp-small-servers
service pt-vty-logging
!
<output truncated>
```

Working with Configuration Files

You can copy (*download*) configuration files from a TFTP, FTP, or RCP server to the running configuration or startup configuration of the switch. You might want to perform this for one of these reasons:

- To restore a backed-up configuration file.
- To use the configuration file for another switch. For example, you might add another switch to your network and want it to have a configuration similar to the original switch. By copying the file to the new switch, you can change the relevant parts rather than recreating the whole file.
- To load the same configuration commands on all the switches in your network so that all the switches have similar configurations.

You can copy (*upload*) configuration files from the switch to a file server by using TFTP, FTP, or RCP. You might perform this task to back up a current configuration file to a server before changing its contents so that you can later restore the original configuration file from the server.

The protocol that you use depends on which type of server you are using. The FTP and RCP transport mechanisms provide faster performance and more reliable delivery of data than TFTP. These improvements are possible because FTP and RCP are built on and use the Transmission Control Protocol/Internet Protocol (TCP/IP) stack, which is connection-oriented.

This section includes this information:

- Guidelines for Creating and Using Configuration Files, on page 474
- Configuration File Types and Location, on page 475
- Creating a Configuration File By Using a Text Editor, on page 475
- Copying Configuration Files By Using TFTP, on page 475
- Copying Configuration Files By Using FTP, on page 477
- Copying Configuration Files By Using RCP, on page 480
- Clearing Configuration Information, on page 484

Guidelines for Creating and Using Configuration Files

Creating configuration files can aid in your switch configuration. Configuration files can contain some or all of the commands needed to configure one or more switches. For example, you might want to download the same configuration file to several switches that have the same hardware configuration.

Use these guidelines when creating a configuration file:

- We recommend that you connect through the service port when using configuration files to configure the switch. If you configure the switch from a Telnet session, IP addresses are not changed, and ports are not disabled.
- If no passwords have been set on the switch, you must set them on each switch by entering the **enable secret** *secret-password* global configuration command. Enter a blank line for this command. The password is saved in the configuration file as clear text.
- If passwords already exist, you cannot enter the **enable secret** *secret-password* global configuration command in the file because the password verification will fail. If you enter a password in the configuration file, the switch mistakenly attempts to execute the passwords as commands as it executes the file.

Note: The **copy {ftp: | rcp: | tftp:} system:running-config** privileged EXEC command loads the configuration files on the switch as if you were entering the commands at the command line. The switch does not erase the existing running configuration before adding the commands. If a command in the copied configuration file replaces a command in the existing configuration file, the existing command is erased. For example, if the copied configuration file contains a different IP address in a particular command than the existing configuration, the IP address in the copied configuration is used. However, some commands in the existing configuration might not be replaced or negated. In this case, the resulting configuration file is a mixture of the existing configuration file and the copied configuration file, with the copied configuration file having precedence.

To restore a configuration file to an exact copy of a file stored on a server, copy the configuration file directly to the startup configuration (by using the **copy {ftp: | rcp: | tftp:} nvram:startup-config** privileged EXEC command), and reload the switch.

Configuration File Types and Location

Startup configuration files are used during system startup to configure the software. Running configuration files contain the current configuration of the software. The two configuration files can be different. For example, you might want to change the configuration for a short time period rather than permanently. In this case, you would change the running configuration but not save the configuration by using the **copy running-config startup-config** privileged EXEC command.

The running configuration is saved in DRAM; the startup configuration is stored in the NVRAM section of flash memory.

Creating a Configuration File By Using a Text Editor

When creating a configuration file, you must list commands logically so that the system can respond appropriately. This is one method of creating a configuration file:

1. Copy an existing configuration from a switch to a server.

For more information, see the “Downloading the Configuration File By Using TFTP” section on page 476, the “Downloading a Configuration File By Using FTP” section on page 478, or the “Downloading a Configuration File By Using RCP” section on page 482.

2. Open the configuration file in a text editor, such as vi or emacs on UNIX or Notepad on a PC.
 3. Extract the portion of the configuration file with the desired commands, and save it in a new file.
 4. Copy the configuration file to the appropriate server location. For example, copy the file to the TFTP directory on the workstation (usually `/tftpboot` on a UNIX workstation).
 5. Make sure the permissions on the file are set to world-read.
-

Copying Configuration Files By Using TFTP

You can configure the switch by using configuration files you create, download from another switch, or download from a TFTP server. You can copy (upload) configuration files to a TFTP server for storage.

This section includes this information:

- Preparing to Download or Upload a Configuration File By Using TFTP, on page 476
- Downloading the Configuration File By Using TFTP, on page 476
- Uploading the Configuration File By Using TFTP, on page 477

Preparing to Download or Upload a Configuration File By Using TFTP

Before you begin downloading or uploading a configuration file by using TFTP, do these tasks:

- Ensure that the workstation acting as the TFTP server is properly configured. On a Sun workstation, make sure that the `/etc/inetd.conf` file contains this line:

```
tftp dgram udp wait root /usr/etc/in.tftpd in.tftpd -p -s /tftpboot
```

Make sure that the `/etc/services` file contains this line:

```
tftp 69/udp
```

Note: You must restart the `inetd` daemon after modifying the `/etc/inetd.conf` and `/etc/services` files. To restart the daemon, either stop the `inetd` process and restart it, or enter a **fastboot** command (on the SunOS 4.x) or a **reboot** command (on Solaris 2.x or SunOS 5.x). For more information on the TFTP daemon, refer to the documentation for your workstation.

- Ensure that the switch has a route to the TFTP server. The switch and the TFTP server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the TFTP server by using the **ping** command.
- Ensure that the configuration file to be downloaded is in the correct directory on the TFTP server (usually `/tftpboot` on a UNIX workstation).
- For download operations, ensure that the permissions on the file are set correctly. The permission on the file should be world-read.
- Before uploading the configuration file, you might need to create an empty file on the TFTP server. To create an empty file, enter the **touch filename** command, where *filename* is the name of the file you will use when uploading it to the server.
- During upload operations, if you are overwriting an existing file (including an empty file, if you had to create one) on the server, ensure that the permissions on the file are set correctly. Permissions on the file should be world-write.

Downloading the Configuration File By Using TFTP

To configure the switch by using a configuration file downloaded from a TFTP server, follow these steps:

1. Copy the configuration file to the appropriate TFTP directory on the workstation.
2. Verify that the TFTP server is properly configured by referring to the “Preparing to Download or Upload a Configuration File By Using TFTP” section on page 476.
3. Log into the switch through a Telnet session.
4. Download the configuration file from the TFTP server to configure the switch.

Specify the IP address or host name of the TFTP server and the name of the file to download.

Use one of these privileged EXEC commands:

- **copy tftp:[*[[[//location]]/directory]]/filename*] system:running-config**

- **copy tftp:[[/location]/directory]/filename] nvram:startup-config**

The configuration file downloads, and the commands are executed as the file is parsed line-by-line.

This example shows how to configure the software from the file *tokyo-config* at IP address 172.16.2.155:

```
Switch# copy tftp://172.16.2.155/tokyo-config system:running-config
Configure using tokyo-config from 172.16.2.155? [confirm] y
Booting tokyo-config from 172.16.2.155:!!! [OK - 874/16000 bytes]
```

Uploading the Configuration File By Using TFTP

To upload a configuration file from a switch to a TFTP server for storage, follow these steps:

1. Verify that the TFTP server is properly configured by referring to the “Preparing to Download or Upload a Configuration File By Using TFTP” section on page 476.
2. Log into the switch through a Telnet session.
3. Upload the switch configuration to the TFTP server. Specify the IP address or host name of the TFTP server and the destination filename.

Use one of these privileged EXEC commands:

- **copy system:running-config tftp:[[/location]/directory]/filename]**
- **copy nvram:startup-config tftp:[[/location]/directory]/filename]**

The file is uploaded to the TFTP server.

This example shows how to upload a configuration file from a switch to a TFTP server:

```
Switch# copy system:running-config tftp://172.16.2.155/tokyo-config
Write file tokyo-config on host 172.16.2.155? [confirm] y
#
Writing tokyo-config!!! [OK]
```

Copying Configuration Files By Using FTP

You can copy configuration files to or from an FTP server.

The FTP protocol requires a client to send a remote username and password on each FTP request to a server. When you copy a configuration file from the switch to a server by using FTP, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **copy** command if a username is specified.
- The username set by the **ip ftp username username** global configuration command if the command is configured.
- Anonymous.

The switch sends the first valid password in this list:

- The password specified in the **copy** command if a password is specified.
- The password set by the **ip ftp password password** global configuration command if the command is configured.
- The switch forms a password named *username@switchname.domain*. The variable *username* is the username associated with the current session, *switchname* is the configured host name, and *domain* is the domain of the switch.

The username and password must be associated with an account on the FTP server. If you are writing to the server, the FTP server must be properly configured to accept your FTP write request.

Use the **ip ftp username** and **ip ftp password** commands to specify a username and password for all copies. Include the username in the **copy** command if you want to specify only a username for that copy operation.

If the server has a directory structure, the configuration file is written to or copied from the directory associated with the username on the server. For example, if the configuration file is in the home directory of a user on the server, specify that user's name as the remote username.

For more information, refer to the documentation for your FTP server.

This section includes this information:

- Preparing to Download or Upload a Configuration File By Using FTP, on page 478
- Downloading a Configuration File By Using FTP, on page 478
- Uploading a Configuration File By Using FTP, on page 479

Preparing to Download or Upload a Configuration File By Using FTP

Before you begin downloading or uploading a configuration file by using FTP, do these tasks:

- Ensure that the switch has a route to the FTP server. The switch and the FTP server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the FTP server by using the **ping** command.
- If you are accessing the switch through a Telnet session and you do not have a valid username, make sure that the current FTP username is the one that you want to use for the FTP download. You can enter the **show users** privileged EXEC command to view the valid username. If you do not want to use this username, create a new FTP username by using the **ip ftp username username** global configuration command during all copy operations. The new username is stored in NVRAM. If you are accessing the switch through a Telnet session and you have a valid username, this username is used, and you do not need to set the FTP username. Include the username in the **copy** command if you want to specify a username for only that copy operation.
- When you upload a configuration file to the FTP server, it must be properly configured to accept the write request from the user on the switch.

For more information, refer to the documentation for your FTP server.

Downloading a Configuration File By Using FTP

Beginning in privileged EXEC mode, follow these steps to download a configuration file by using FTP:

Step	Command	Purpose
1.		Verify that the FTP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using FTP" section on page 478.
2.		Log into the switch through a Telnet session.

Step	Command	Purpose
3.	configure terminal	Enter global configuration mode on the switch. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
4.	ip ftp username <i>username</i>	(Optional) Change the default remote username.
5.	ip ftp password <i>password</i>	(Optional) Change the default password.
6.	end	Return to privileged EXEC mode.
7.	copy ftp:[[[/[username[:password]@]/location]/directory]/filename] system:running-config or copy ftp:[[[/[username[:password]@]/location]/directory]/filename] nvrnram:startup-config	Using FTP, copy the configuration file from a network server to the running configuration or to the startup configuration file.

This example shows how to copy a configuration file named *host1-config* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 and to load and run those commands on the switch:

```
Switch# copy ftp://netadmin1:mypass@172.16.101.101/host1-config
system:running-config
Configure using host1-config from 172.16.101.101? [confirm]
Connected to 172.16.101.101
Loading 1112 byte file host1-config:[OK]
Switch#
%SYS-5-CONFIG: Configured from host1-config by ftp from 172.16.101.101
```

This example shows how to specify a remote username of *netadmin1*. The software copies the configuration file *host2-config* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 to the switch startup configuration.

```
Switch# configure terminal
Switch(config)# ip ftp username netadmin1
Switch(config)# ip ftp password mypass
Switch(config)# end
Switch# copy ftp: nvrnram:startup-config
Address of remote host [255.255.255.255]? 172.16.101.101
Name of configuration file[rtr2-config]? host2-config
Configure using host2-config from 172.16.101.101?[confirm]
Connected to 172.16.101.101
Loading 1112 byte file host2-config:[OK]
[OK]
Switch#
%SYS-5-CONFIG_NV:Non-volatile store configured from host2-config by ftp from
172.16.101.101
```

Uploading a Configuration File By Using FTP

Beginning in privileged EXEC mode, follow these steps to upload a configuration file by using FTP:

Step	Command	Purpose
1.		Verify that the FTP server is properly configured by referring to the “Preparing to Download or Upload a Configuration File By Using FTP” section on page 478.
2.		Log into the switch through a Telnet session.
3.	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
4.	ip ftp username <i>username</i>	(Optional) Change the default remote username.
5.	ip ftp password <i>password</i>	(Optional) Change the default password.
6.	end	Return to privileged EXEC mode.
7.	copy system:running-config ftp:[[[/[username[:password]@]/location]/directory]/filename] or copy nvram:startup-config ftp:[[[/[username[:password]@]/location]/directory]/filename]	Using FTP, store the switch running or startup configuration file to the specified location.

This example shows how to copy the running configuration file named *switch2-config* to the *netadmin1* directory on the remote host with an IP address of 172.16.101.101:

```
Switch# copy system:running-config
ftp://netadmin1:mypass@172.16.101.101/switch2-config
Write file switch2-config on host 172.16.101.101?[confirm]
Building configuration...[OK]
Connected to 172.16.101.101
Switch#
```

This example shows how to store a startup configuration file on a server by using FTP to copy the file:

```
Switch# configure terminal
Switch(config)# ip ftp username netadmin2
Switch(config)# ip ftp password mypass
Switch(config)# end
Switch# copy nvram:startup-config ftp:
Remote host[]? 172.16.101.101
Name of configuration file to write [switch2-config]?
Write file switch2-config on host 172.16.101.101?[confirm]
![OK]
```

Copying Configuration Files By Using RCP

The Remote Copy Protocol (RCP) provides another method of downloading, uploading, and copying configuration files between remote hosts and the switch. Unlike TFTP, which uses User Datagram Protocol (UDP), a connectionless protocol, RCP uses TCP, which is connection-oriented.

To use RCP to copy files, the server from or to which you will be copying files must support RCP. The RCP copy commands rely on the rsh server (or daemon) on the

remote system. To copy files by using RCP, you do not need to create a server for file distribution as you do with TFTP. You only need to have access to a server that supports the remote shell (rsh). (Most UNIX systems support rsh.) Because you are copying a file from one place to another, you must have read permission on the source file and write permission on the destination file. If the destination file does not exist, RCP creates it for you.

The RCP requires a client to send a remote username with each RCP request to a server. When you copy a configuration file from the switch to a server, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **copy** command if a username is specified.
- The username set by the **ip rcmd remote-username username** global configuration command if the command is configured.
- The remote username associated with the current TTY (terminal) process. For example, if the user is connected to the router through Telnet and was authenticated through the **username** command, the switch software sends the Telnet username as the remote username.
- The switch host name.

For a successful RCP copy request, you must define an account on the network server for the remote username. If the server has a directory structure, the configuration file is written to or copied from the directory associated with the remote username on the server. For example, if the configuration file is in the home directory of a user on the server, specify that user's name as the remote username.

This section includes this information:

- Preparing to Download or Upload a Configuration File By Using RCP, on page 481
- Downloading a Configuration File By Using RCP, on page 482
- Uploading a Configuration File By Using RCP, on page 483

Preparing to Download or Upload a Configuration File By Using RCP

Before you begin downloading or uploading a configuration file by using RCP, do these tasks:

- Ensure that the workstation acting as the RCP server supports the remote shell (rsh).
- Ensure that the switch has a route to the RCP server. The switch and the server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the RCP server by using the **ping** command.
- If you are accessing the switch through a Telnet session and you do not have a valid username, make sure that the current RCP username is the one that you want to use for the RCP download. You can enter the **show users** privileged EXEC command to view the valid username. If you do not want to use this username, create a new RCP username by using the **ip rcmd remote-username username** global configuration command to be used during all copy operations. The new username is stored in NVRAM. If you are accessing the switch through a Telnet session and you have a valid username, this username is used, and you do not need to set the RCP username. Include the username in the **copy** command if you want to specify a username for only that copy operation.
- When you upload a file to the RCP server, it must be properly configured to accept the RCP write request from the user on the switch. For UNIX systems, you must

add an entry to the .rhosts file for the remote user on the RCP server. For example, suppose that the switch contains these configuration lines:

```
hostname Switch1
ip rcmd remote-username User0
```

If the switch IP address translates to *Switch1.company.com*, the .rhosts file for User0 on the RCP server should contain this line:

```
Switch1.company.com Switch1
```

For more information, refer to the documentation for your RCP server.

Downloading a Configuration File By Using RCP

Beginning in privileged EXEC mode, follow these steps to download a configuration file by using RCP:

Step	Command	Purpose
1.		Verify that the RCP server is properly configured by referring to the “Preparing to Download or Upload a Configuration File By Using RCP” section on page 481.
2.		Log into the switch through a Telnet session.
3.	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
4.	ip rcmd remote-username <i>username</i>	(Optional) Specify the remote username.
5.	end	Return to privileged EXEC mode.
6.	copy rcp:[[[/[username@]location]/directory]/filename] system:running-config or copy rcp:[[[/[username@]location]/directory]/filename] nvram:startup-config	Using RCP, copy the configuration file from a network server to the running configuration or to the startup configuration file.

This example shows how to copy a configuration file named *host1-config* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 and load and run those commands on the switch:

```
Switch# copy rcp://netadmin1@172.16.101.101/host1-config system:running-
config
Configure using host1-config from 172.16.101.101? [confirm]
Connected to 172.16.101.101
Loading 1112 byte file host1-config:![OK]
Switch#
%SYS-5-CONFIG: Configured from host1-config by rcp from 172.16.101.101
```

This example shows how to specify a remote username of *netadmin1*. Then it copies the configuration file *host2-confg* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 to the startup configuration:

```
Switch# configure terminal
Switch(config)# ip rcmd remote-username netadmin1
Switch(config)# end
Switch# copy rcp: nvram:startup-config
Address of remote host [255.255.255.255]? 172.16.101.101
Name of configuration file[rtr2-confg]? host2-confg
Configure using host2-confg from 172.16.101.101?[confirm]
Connected to 172.16.101.101
Loading 1112 byte file host2-confg:![OK]
[OK]
Switch#
%SYS-5-CONFIG_NV:Non-volatile store configured from host2-confg by rcp from
172.16.101.101
```

Uploading a Configuration File By Using RCP

Beginning in privileged EXEC mode, follow these steps to upload a configuration file by using RCP:

Step	Command	Purpose
1.		Verify that the RCP server is properly configured by referring to the “Preparing to Download or Upload a Configuration File By Using RCP” section on page 481.
2.		Log into the switch through a Telnet session.
3.	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
4.	ip rcmd remote-username <i>username</i>	(Optional) Specify the remote username.
5.	end	Return to privileged EXEC mode.
6.	copy system:running-config rcp:[[[/[username@]location]/directory]/filename] or copy nvram:startup-config rcp:[[[/[username@]location]/directory]/filename]	Using RCP, copy the configuration file from a switch running or startup configuration file to a network server.

This example shows how to copy the running configuration file named *switch2-confg* to the *netadmin1* directory on the remote host with an IP address of 172.16.101.101:

```
Switch# copy system:running-config rcp://netadmin1@172.16.101.101/switch2-  
confg
Write file switch-confg on host 172.16.101.101?[confirm]
Building configuration...[OK]
Connected to 172.16.101.101
Switch#
```

This example shows how to store a startup configuration file on a server:

```
Switch# configure terminal
Switch(config)# ip rcmd remote-username netadmin2
Switch(config)# end
Switch# copy nvram:startup-config rcp:
Remote host[]? 172.16.101.101
Name of configuration file to write [switch2-config]?
Write file switch2-config on host 172.16.101.101?[confirm]
![OK]
```

Clearing Configuration Information

You can clear the configuration information from the startup configuration.

Clearing the Startup Configuration File

To clear the contents of your startup configuration, use the **erase nvram:** or the **erase startup-config** privileged EXEC command.

Caution: You cannot restore the startup configuration file after it has been deleted.

Deleting a Stored Configuration File

To delete a saved configuration from flash memory, use the **delete flash:filename** privileged EXEC command. Depending on the setting of the **file prompt** global configuration command, you might be prompted for confirmation before you delete a file. By default, the switch prompts for confirmation on destructive file operations. For more information about the **file prompt** command, refer to the *Cisco IOS Command Reference for Cisco IOS Release 12.1*.

Caution: You cannot restore a file after it has been deleted.

Working with Software Images

This section describes how to archive (download and upload) software image files, which contain the system software, Cisco IOS code, and the web management HTML files.

You download a switch image file from a TFTP, FTP, or RCP server to upgrade the switch software. You can replace the current image with the new one or keep the current image in flash memory after a download.

You upload a switch image file to a TFTP, FTP, or RCP server for backup purposes. You can use this uploaded image for future downloads to the same switch or another of the same type.

The protocol that you use depends on which type of server that you are using. The FTP and RCP transport mechanisms provide faster performance and more reliable delivery of data than TFTP. These improvements are possible because FTP and RCP are built on and use the Transmission Control Protocol/Internet Protocol (TCP/IP) stack, which is connection-oriented.

This section includes this information:

- Image Location on the Switch, on page 485
- tar File Format of Images on a Server or Cisco.com, on page 485
- Copying Image Files By Using TFTP, on page 486
- Copying Image Files By Using FTP, on page 489

- Copying Image Files By Using RCP, on page 493

Note: For a list of software images and the supported upgrade paths, refer to the release notes.

Image Location on the Switch

The Cisco IOS image is stored as a *.bin* file in a directory that shows the version number. A subdirectory contains the HTML files needed for web management. The image is stored on the system board flash memory (flash:).

You can use the **show version** privileged EXEC command to see the software version that is currently running on your switch. In the display, check the line that begins with `System image file is...`. It shows the directory name in flash memory where the image is stored.

You can also use the **dir filesystem:** privileged EXEC command to see the directory names of other software images you might have stored in flash memory.

tar File Format of Images on a Server or Cisco.com

Software images located on a server or downloaded from Cisco.com are provided in a tar file format, which contains these files:

- *info* file
The info file is always at the beginning of the tar file and has information about the files within it.
- Cisco IOS image
- Web management files needed by the HTTP server on the switch
- *info.ver* file
The info.ver file is always at the end of the tar file and has the same information as the info file. Because it is the last file in the tar file, its existence means that all files in the image have been downloaded.

This example shows the information in the info and info.ver files:

```
version_suffix: i6q412-121-0.0.42.EA1
version_directory: cigesm-i6q412-mz.121-0.0.42.EA1
image_name: cigesm-i6q412-mz.121-0.0.42.EA1.bin
ios_image_file_size:3038720
total_image_file_size: 5404672
image_feature: LAYER_2|MIN_DRAM_MEG=32
image_family: CIGESM
image_min_dram: 32
info_end:
```

Table 97. *info* and *info.ver* File Description

Field	Description
version_suffix	Specifies the Cisco IOS image version string suffix
version_directory	Specifies the directory where the Cisco IOS image and the HTML subdirectory are installed
image_name	Specifies the name of the Cisco IOS image within the tar file
ios_image_file_size	Specifies the Cisco IOS image size in the tar file, which is an approximate measure of how much flash space is required to hold just the Cisco IOS image

Table 97. *info* and *info.ver* File Description (continued)

Field	Description
total_image_file_size	Specifies the size of all the images (the Cisco IOS image and the HTML files) in the tar file, which is an approximate measure of how much flash space is required to hold them
image_feature	Describes the core functionality of the image
image_family	Describes the family of products on which the software can be installed
image_min_dram	Specifies the minimum amount of DRAM needed to run this image

Copying Image Files By Using TFTP

You can download a switch image from a TFTP server or upload the image from the switch to a TFTP server.

You download a switch image file from a server to upgrade the switch software. You can overwrite the current image with the new one or keep the current image after a download.

You upload a switch image file to a server for backup purposes; this uploaded image can be used for future downloads to the same or another switch of the same type.

This section includes this information:

- Preparing to Download or Upload an Image File By Using TFTP, on page 486
- Downloading an Image File By Using TFTP, on page 487
- Uploading an Image File By Using TFTP, on page 488

Preparing to Download or Upload an Image File By Using TFTP

Before you begin downloading or uploading an image file by using TFTP, do these tasks:

- Ensure that the workstation acting as the TFTP server is properly configured. On a Sun workstation, make sure that the `/etc/inetd.conf` file contains this line:

```
tftp dgram udp wait root /usr/etc/in.tftpd in.tftpd -p -s /tftpboot
```

Make sure that the `/etc/services` file contains this line:

```
tftp 69/udp
```

Note: You must restart the `inetd` daemon after modifying the `/etc/inetd.conf` and `/etc/services` files. To restart the daemon, either stop the `inetd` process and restart it, or enter a **fastboot** command (on the SunOS 4.x) or a **reboot** command (on Solaris 2.x or SunOS 5.x). For more information on the TFTP daemon, refer to the documentation for your workstation.

- Ensure that the switch has a route to the TFTP server. The switch and the TFTP server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the TFTP server by using the **ping** command.
- Ensure that the image to be downloaded is in the correct directory on the TFTP server (usually `/tftpboot` on a UNIX workstation).
- For download operations, ensure that the permissions on the file are set correctly. The permission on the file should be world-read.
- Before uploading the image file, you might need to create an empty file on the TFTP server. To create an empty file, enter the **touch filename** command, where

filename is the name of the file you will use when uploading the image to the server.

- During upload operations, if you are overwriting an existing file (including an empty file, if you had to create one) on the server, ensure that the permissions on the file are set correctly. Permissions on the file should be world-write.

Downloading an Image File By Using TFTP

You can download a new image file and replace the current image or keep the current image.

Beginning in privileged EXEC mode, follow Steps 1 through 3 to download a new image from a TFTP server and overwrite the existing image. To keep the current image, skip Step 3.

Step	Command	Purpose
1.		Copy the image to the appropriate TFTP directory on the workstation. Make sure that the TFTP server is properly configured; see the “Preparing to Download or Upload an Image File By Using TFTP” section on page 486.
2.		Log into the switch through a Telnet session.
3.	archive download-sw /overwrite /reload tftp:[<i>///location</i>]/<i>directory</i>]/<i>image-name.tar</i>	<p>Download the image file from the TFTP server to the switch, and overwrite the current image.</p> <ul style="list-style-type: none"> • The /overwrite option overwrites the software image in flash with the downloaded image only if the version of the image being downloaded is the same as the existing copy in flash memory. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For <i>///location</i>, specify the IP address of the TFTP server. • For <i>/directory/image-name.tar</i>, specify the directory (optional) and the image to download. Directory and image names are case sensitive.
4.	archive download-sw /leave-old-sw /reload tftp:[<i>///location</i>]/<i>directory</i>]/<i>image-name.tar</i>	<p>Download the image file from the TFTP server to the switch, and keep the current image.</p> <ul style="list-style-type: none"> • The /leave-old-sw option keeps the old software version after a download. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For <i>///location</i>, specify the IP address of the TFTP server. • For <i>/directory/image-name.tar</i>, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it stops the process and reports an error. If you specify the **/overwrite** option, the download algorithm removes the existing image on the flash device whether or not it is the same as the new one, downloads the new image, and then reloads the software.

Note: If the flash device has sufficient space to hold two images and you want to overwrite one of these images with the same version, you must specify the **/overwrite** option.

If you specify the **/leave-old-sw**, the existing files are not removed. If there is not enough space to install the new image and keep the current running image, the download process stops, and an error message appears.

The algorithm installs the downloaded image on the system board flash device (flash:). The image is placed into a new directory named with the software version string, and the BOOT environment variable is updated to point to the newly installed image.

If you kept the old image during the download process (you specified the **/leave-old-sw** keyword), you can remove it by entering the **delete /force /recursive filesystem:/file-url** privileged EXEC command. For *filesystem*, use **flash:** for the system board flash device. For *file-url*, enter the directory name of the old image. All the files in the directory and the directory are removed.

Caution: For the download and upload algorithms to operate properly, do *not* rename image names.

Uploading an Image File By Using TFTP

You can upload an image from the switch to a TFTP server. You can later download this image to the switch or to another switch of the same type.

The upload feature is available only if the HTML pages associated with the Cluster Management Suite (CMS) have been installed with the existing image.

Beginning in privileged EXEC mode, follow these steps to upload an image to a TFTP server:

Step	Command	Purpose
1.		Make sure that the TFTP server is properly configured; see the “Preparing to Download or Upload an Image File By Using TFTP” section on page 486.
1.		Log into the switch through a Telnet session.
2.	archive upload-sw tftp:[[/location]/directory]image-name.tar	Upload the currently running switch image to the TFTP server. <ul style="list-style-type: none"> For <i>[/location]</i>, specify the IP address of the TFTP server. For <i>/directory/image-name.tar</i>, specify the directory (optional) and the name of the software image to be uploaded. Directory and image names are case sensitive. The <i>image-name.tar</i> is the name of the software image to be stored on the server.

The **archive upload-sw** privileged EXEC command builds an image file on the server by uploading these files in order: info, the Cisco IOS image, the HTML files, and info.ver. After these files are uploaded, the upload algorithm creates the tar file format.

Caution: For the download and upload algorithms to operate properly, do *not* rename image names.

Copying Image Files By Using FTP

You can download a switch image from an FTP server or upload the image from the switch to an FTP server.

You download a switch image file from a server to upgrade the switch software. You can overwrite the current image with the new one or keep the current image after a download.

You upload a switch image file to a server for backup purposes. You can use this uploaded image for future downloads to the switch or another switch of the same type.

This section includes this information:

- Preparing to Download or Upload an Image File By Using FTP, on page 489
- Downloading an Image File By Using FTP, on page 490
- Uploading an Image File By Using FTP, on page 492

Preparing to Download or Upload an Image File By Using FTP

You can copy images files to or from an FTP server.

The FTP protocol requires a client to send a remote username and password on each FTP request to a server. When you copy an image file from the switch to a server by using FTP, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **archive download-sw** or **archive upload-sw** privileged EXEC command if a username is specified.
- The username set by the **ip ftp username** *username* global configuration command if the command is configured.
- Anonymous.

The switch sends the first valid password in this list:

- The password specified in the **archive download-sw** or **archive upload-sw** privileged EXEC command if a password is specified.
- The password set by the **ip ftp password** *password* global configuration command if the command is configured.
- The switch forms a password named *username@switchname.domain*. The variable *username* is the username associated with the current session, *switchname* is the configured host name, and *domain* is the domain of the switch.

The username and password must be associated with an account on the FTP server. If you are writing to the server, the FTP server must be properly configured to accept the FTP write request from you.

Use the **ip ftp username** and **ip ftp password** commands to specify a username and password for all copies. Include the username in the **archive download-sw** or **archive upload-sw** privileged EXEC command if you want to specify a username only for that operation.

If the server has a directory structure, the image file is written to or copied from the directory associated with the username on the server. For example, if the image file resides in the home directory of a user on the server, specify that user's name as the remote username.

Before you begin downloading or uploading an image file by using FTP, do these tasks:

- Ensure that the switch has a route to the FTP server. The switch and the FTP server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the FTP server by using the **ping** command.
- If you are accessing the switch through a Telnet session and you do not have a valid username, make sure that the current FTP username is the one that you want to use for the FTP download. You can enter the **show users** privileged EXEC command to view the valid username. If you do not want to use this username, create a new FTP username by using the **ip ftp username *username*** global configuration command. This new name will be used during all archive operations. The new username is stored in NVRAM. If you are accessing the switch through a Telnet session and you have a valid username, this username is used, and you do not need to set the FTP username. Include the username in the **archive download-sw** or **archive upload-sw** privileged EXEC command if you want to specify a username for that operation only.
- When you upload an image file to the FTP server, it must be properly configured to accept the write request from the user on the switch.

For more information, refer to the documentation for your FTP server.

Downloading an Image File By Using FTP

You can download a new image file and overwrite the current image or keep the current image.

Beginning in privileged EXEC mode, follow Steps 1 through 7 to download a new image from an FTP server and overwrite the existing image. To keep the current image, skip Step 7.

Step	Command	Purpose
1.		Verify that the FTP server is properly configured by referring to the “Preparing to Download or Upload an Image File By Using FTP” section on page 489.
2.		Log into the switch through a Telnet session.
3.	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
4.	ip ftp username <i>username</i>	(Optional) Change the default remote username.
5.	ip ftp password <i>password</i>	(Optional) Change the default password.
6.	end	Return to privileged EXEC mode.

Step	Command	Purpose
7.	archive download-sw /overwrite /reload ftp:[//username[:password]@location]/directory/image-name.tar	Download the image file from the FTP server to the switch, and overwrite the current image. <ul style="list-style-type: none"> • The /overwrite option overwrites the software image in flash with the downloaded image. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For //username[:password], specify the username and password; these must be associated with an account on the FTP server. For more information, see the “Preparing to Download or Upload an Image File By Using FTP” section on page 489. • For @location, specify the IP address of the FTP server. • For directory/image-name.tar, specify the directory (optional) and the image to download. Directory and image names are case sensitive.
8.	archive download-sw /leave-old-sw /reload ftp:[//username[:password]@location]/directory/image-name.tar	Download the image file from the FTP server to the switch, and keep the current image. <ul style="list-style-type: none"> • The /leave-old-sw option keeps the old software version after a download. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For //username[:password], specify the username and password. These must be associated with an account on the FTP server. For more information, see the “Preparing to Download or Upload an Image File By Using FTP” section on page 489. • For @location, specify the IP address of the FTP server. • For directory/image-name.tar, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it stops the process and reports an error. If you specify the **/overwrite** option, the download algorithm removes the existing image on the flash device, whether or not it is the same as the new one, downloads the new image, and then reloads the software.

Note: If the flash device has sufficient space to hold two images and you want to overwrite one of these images with the same version, you must specify the **/overwrite** option.

If you specify the **/leave-old-sw**, the existing files are not removed. If there is not enough space to install the new image and keep the running image, the download process stops, and an error message appears.

The algorithm installs the downloaded image onto the system board flash device (flash:). The image is placed into a new directory named with the software version

string, and the BOOT environment variable is updated to point to the newly installed image.

If you kept the old image during the download process (you specified the **/leave-old-sw** keyword), you can remove it by entering the **delete /force /recursive filesystem:/file-url** privileged EXEC command. For *filesystem*, use **flash:** for the system board flash device. For *file-url*, enter the directory name of the old software image. All the files in the directory and the directory are removed.

Caution: For the download and upload algorithms to operate properly, do *not* rename image names.

Uploading an Image File By Using FTP

You can upload an image from the switch to an FTP server. You can later download this image to the same switch or to another switch of the same type.

The upload feature is available only if the HTML pages associated with the Cluster Management Suite (CMS) have been installed with the existing image.

Beginning in privileged EXEC mode, follow these steps to upload an image to an FTP server:

Step	Command	Purpose
1.		Verify that the FTP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using FTP" section on page 478.
2.		Log into the switch through a Telnet session.
3.	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
4.	ip ftp username <i>username</i>	(Optional) Change the default remote username.
5.	ip ftp password <i>password</i>	(Optional) Change the default password.
6.	end	Return to privileged EXEC mode.
7.	archive upload-sw ftp:[//[username[:password]@]location/]directory/]image-name.tar	Upload the currently running switch image to the FTP server. <ul style="list-style-type: none"> For <i>//username:password</i>, specify the username and password. These must be associated with an account on the FTP server. For more information, see the "Preparing to Download or Upload an Image File By Using FTP" section on page 489. For <i>@location</i>, specify the IP address of the FTP server. For <i>/directory/image-name.tar</i>, specify the directory (optional) and the name of the software image to be uploaded. Directory and image names are case sensitive. The <i>image-name.tar</i> is the name of the software image to be stored on the server.

The **archive upload-sw** command builds an image file on the server by uploading these files in order: info, the Cisco IOS image, the HTML files, and info.ver. After these files are uploaded, the upload algorithm creates the tar file format.

Caution: For the download and upload algorithms to operate properly, do *not* rename image names.

Copying Image Files By Using RCP

You can download a switch image from an RCP server or upload the image from the switch to an RCP server.

You download a switch image file from a server to upgrade the switch software. You can overwrite the current image with the new one or keep the current image after a download.

You upload a switch image file to a server for backup purposes. You can use this uploaded image for future downloads to the same switch or another of the same type.

This section includes this information:

- Preparing to Download or Upload an Image File By Using RCP, on page 493
- Downloading an Image File By Using RCP, on page 494
- Uploading an Image File By Using RCP, on page 496

Preparing to Download or Upload an Image File By Using RCP

RCP provides another method of downloading and uploading image files between remote hosts and the switch. Unlike TFTP, which uses User Datagram Protocol (UDP), a connectionless protocol, RCP uses TCP, which is connection-oriented.

To use RCP to copy files, the server from or to which you will be copying files must support RCP. The RCP copy commands rely on the rsh server (or daemon) on the remote system. To copy files by using RCP, you do not need to create a server for file distribution as you do with TFTP. You only need to have access to a server that supports the remote shell (rsh). (Most UNIX systems support rsh.) Because you are copying a file from one place to another, you must have read permission on the source file and write permission on the destination file. If the destination file does not exist, RCP creates it for you.

RCP requires a client to send a remote username on each RCP request to a server. When you copy an image from the switch to a server by using RCP, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **archive download-sw** or **archive upload-sw** privileged EXEC command if a username is specified.
- The username set by the **ip rcmd remote-username *username*** global configuration command if the command is entered.
- The remote username associated with the current TTY (terminal) process. For example, if the user is connected to the router through Telnet and was authenticated through the **username** command, the switch software sends the Telnet username as the remote username.
- The switch host name.

For the RCP copy request to execute successfully, an account must be defined on the network server for the remote username. If the server has a directory structure, the image file is written to or copied from the directory associated with the remote username on the server. For example, if the image file resides in the home directory of a user on the server, specify that user's name as the remote username.

Before you begin downloading or uploading an image file by using RCP, do these tasks:

- Ensure that the workstation acting as the RCP server supports the remote shell (rsh).
- Ensure that the switch has a route to the RCP server. The switch and the server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the RCP server by using the **ping** command.
- If you are accessing the switch through a Telnet session and you do not have a valid username, make sure that the current RCP username is the one that you want to use for the RCP download. You can enter the **show users** privileged EXEC command to view the valid username. If you do not want to use this username, create a new RCP username by using the **ip rcmd remote-username username** global configuration command to be used during all archive operations. The new username is stored in NVRAM. If you are accessing the switch through a Telnet session and you have a valid username, this username is used, and there is no need to set the RCP username. Include the username in the **archive download-sw** or **archive upload-sw** privileged EXEC command if you want to specify a username only for that operation.
- When you upload an image to the RCP to the server, it must be properly configured to accept the RCP write request from the user on the switch. For UNIX systems, you must add an entry to the .rhosts file for the remote user on the RCP server. For example, suppose the switch contains these configuration lines:

```
hostname Switch1
ip rcmd remote-username User0
```

If the switch IP address translates to *Switch1.company.com*, the .rhosts file for User0 on the RCP server should contain this line:

```
Switch1.company.com Switch1
```

For more information, refer to the documentation for your RCP server.

Downloading an Image File By Using RCP

You can download a new image file and replace or keep the current image.

Beginning in privileged EXEC mode, follow Steps 1 through 6 to download a new image from an RCP server and overwrite the existing image. To keep the current image, skip Step 6.

Step	Command	Purpose
1.		Verify that the RCP server is properly configured by referring to the "Preparing to Download or Upload an Image File By Using RCP" section on page 493.
2.		Log into the switch through a Telnet session.
3.	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
4.	ip rcmd remote-username username	(Optional) Specify the remote username.
5.	end	Return to privileged EXEC mode.

Step	Command	Purpose
6.	archive download-sw /overwrite /reload rcp:[[[//[username@]location]/directory]/image-name.tar]	<p>Download the image file from the RCP server to the switch, and overwrite the current image.</p> <ul style="list-style-type: none"> • The /overwrite option overwrites the software image in flash with the downloaded image. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For //username, specify the username. For the RCP copy request to execute successfully, an account must be defined on the network server for the remote username. For more information, see the “Preparing to Download or Upload an Image File By Using RCP” section on page 493. • For @location, specify the IP address of the RCP server. • For /directory/image-name.tar, specify the directory (optional) and the image to download. Directory and image names are case sensitive.
7.	archive download-sw /leave-old-sw /reload rcp:[[[//[username@]location]/directory]/image-name.tar]	<p>Download the image file from the RCP server to the switch, and keep the current image.</p> <ul style="list-style-type: none"> • The /leave-old-sw option keeps the old software version after a download. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For //username, specify the username. For the RCP copy request to execute, an account must be defined on the network server for the remote username. For more information, see the “Preparing to Download or Upload an Image File By Using RCP” section on page 493. • For @location, specify the IP address of the RCP server. • For /directory/image-name.tar, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it stops the process and reports an error. If you specify the **/overwrite** option, the download algorithm removes the existing image on the flash device whether or not it is the same as the new one, downloads the new image, and then reloads the software.

Note: If the flash device has sufficient space to hold two images and you want to overwrite one of these images with the same version, you must specify the **/overwrite** option.

If you specify the **/leave-old-sw**, the existing files are not removed. If there is not enough room to install the new image and keep the running image, the download process stops, and an error message appears.

The algorithm installs the downloaded image onto the system board flash device (flash:). The image is placed in a new directory named with the software version string, and the BOOT environment variable is updated to point to the newly installed image.

If you kept the old software during the download process (you specified the **//leave-old-sw** keyword), you can remove it by entering the **delete //force //recursive filesystem:file-url** privileged EXEC command. For *filesystem*, use **flash:** for the system board flash device. For *file-url*, enter the directory name of the old software image. All the files in the directory and the directory are removed.

Caution: For the download and upload algorithms to operate properly, do *not* rename image names.

Uploading an Image File By Using RCP

You can upload an image from the switch to an RCP server. You can later download this image to the same switch or to another switch of the same type.

The upload feature is available only if the HTML pages associated with the Cluster Management Suite (CMS) have been installed with the existing image.

Beginning in privileged EXEC mode, follow these steps to upload an image to an RCP server:

Step	Command	Purpose
1.		Verify that the RCP server is properly configured by referring to the “Preparing to Download or Upload an Image File By Using RCP” section on page 493.
2.		Log into the switch through a Telnet session.
3.	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
4.	ip rcmd remote-username <i>username</i>	(Optional) Specify the remote username.
5.	end	Return to privileged EXEC mode.
6.	archive upload-sw rcp:[[[//[username@]location]/directory]/image-name.tar]	Upload the currently running switch image to the RCP server. <ul style="list-style-type: none"> For <i>//username</i>, specify the username; for the RCP copy request to execute, an account must be defined on the network server for the remote username. For more information, see the “Preparing to Download or Upload an Image File By Using RCP” section on page 493. For <i>@location</i>, specify the IP address of the RCP server. For <i>/directory]/image-name.tar</i>, specify the directory (optional) and the name of the software image to be uploaded. Directory and image names are case sensitive. The <i>image-name.tar</i> is the name of software image to be stored on the server.

The **archive upload-sw** privileged EXEC command builds an image file on the server by uploading these files in order: info, the Cisco IOS image, the HTML files, and info.ver. After these files are uploaded, the upload algorithm creates the tar file format.

Caution: For the download and upload algorithms to operate properly, do *not* rename image names.

Appendix C. Getting help and technical assistance

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This appendix contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your BladeCenter system, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure that you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure that they are connected.
- Check the power switches to make sure that the system is turned on.
- Use the troubleshooting information in your system documentation, and use the diagnostic tools that come with your system. Information about diagnostic tools is in the *Hardware Maintenance Manual and Troubleshooting Guide* on the IBM *BladeCenter Documentation* CD or at the IBM Support Web site.
- Go to the IBM Support Web site at <http://www.ibm.com/pc/support/> to check for technical information, hints, tips, and new device drivers.

You can solve many problems without outside assistance by following the troubleshooting procedures that IBM provides in the online help or in the publications that are provided with your system and software. The information that comes with your system also describes the diagnostic tests that you can perform. Most xSeries and IntelliStation® systems, operating systems, and programs come with information that contains troubleshooting procedures and explanations of error messages and error codes. If you suspect a software problem, see the information for the operating system or program.

Using the documentation

Information about your IBM BladeCenter, xSeries, or IntelliStation system and preinstalled software, if any, is available in the documentation that comes with your system. That documentation includes printed books, online books, readme files, and help files. See the troubleshooting information in your system documentation for instructions for using the diagnostic programs. The troubleshooting information or the diagnostic programs might tell you that you need additional or updated device drivers or other software. IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. To access these pages, go to <http://www.ibm.com/pc/support/> and follow the instructions. Also, you can order publications through the IBM Publications Ordering System at <http://www.elink.ibm.com/public/applications/publications/cgibin/pbi.cgi>.

Getting help and information from the World Wide Web

On the World Wide Web, the IBM Web site has up-to-date information about IBM BladeCenter, xSeries, and IntelliStation products, services, and support. The address for IBM BladeCenter and xSeries information is <http://www.ibm.com/eserver/xseries/>. The address for IBM IntelliStation information is <http://www.ibm.com/pc/intellistation/>.

You can find service information for your IBM products, including supported options, at <http://www.ibm.com/pc/support/>.

Software service and support

Through IBM Support Line, you can get telephone assistance, for a fee, with usage, configuration, and software problems with BladeCenter and xSeries servers, IntelliStation workstations, and appliances. For information about which products are supported by Support Line in your country or region, go to <http://www.ibm.com/services/sl/products/>.

For more information about Support Line and other IBM services, go to <http://www.ibm.com/services/>, or go to <http://www.ibm.com/planetwide/> for support telephone numbers. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services or through your IBM reseller, if your reseller is authorized by IBM to provide warranty service. Go to <http://www.ibm.com/planetwide/> for support telephone numbers, or in the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

In the U.S. and Canada, hardware service and support is available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9 a.m. to 6 p.m.

Appendix D. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product, and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Edition notice

© Copyright International Business Machines Corporation 2004. All rights reserved.

U.S. Government Users Restricted Rights — Use, duplication, or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

Active Memory	Predictive Failure Analysis
Active PCI	PS/2
Active PCI-X	ServeRAID
Alert on LAN	ServerGuide
BladeCenter	ServerProven
C2T Interconnect	TechConnect
Chipkill	ThinkPad
EtherJet	Tivoli
e-business logo	Tivoli Enterprise
Eserver	Update Connector
FlashCopy	Wake on LAN
IBM	XA-32
IBM (logo)	XA-64
IntelliStation	X-Architecture
NetBAY	XceL4
Netfinity	XpandOnDemand
NetView	xSeries
OS/2 WARP	

Cisco, Cisco IOS, Cisco Systems, the Cisco Systems logo, Catalyst, EtherChannel, IOS, IP/TV, Packet, and SwitchProbe are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

Intel, MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Index

Numerics

- 802.1D
 - See STP
- 802.1Q
 - and trunk ports 166
 - configuration limitations 257
 - native VLAN for untagged traffic 262
 - trunk mode 42
- 802.1S
 - See MSTP
- 802.1W
 - See RSTP
- 802.1X
 - See port-based authentication
- 802.3X flow control 176

A

- abbreviating commands 28
- AC (command switch) 79, 88
- access control entries
 - See ACEs
- access groups, viewing 390
- access lists
 - See ACLs
- access ports
 - defined 165
 - in switch clusters 78
- access-denied response, VMPS 266
- accessing
 - clusters, switch 82
 - command switches 80
 - member switches 82
 - switch clusters 82
- accounting
 - with RADIUS 142
 - with TACACS+ 126, 131
- ACEs
 - defined 372
 - Ethernet 372
 - IP 372
 - Layer 3 parameters 379
 - Layer 4 parameters 379
- ACLs
 - ACEs 372
 - any keyword 378
 - applying
 - time ranges to 383
 - to management interfaces 387
 - to management VLANs 388
 - to physical interfaces 388
 - to QoS 414
 - to terminal lines 387
 - classifying traffic for QoS 428
 - comments in 385
 - compiling 390

- configuration guidelines
 - management interfaces, applying to 375
 - physical interfaces, applying to 375
 - defined 371
 - displaying interface 390
 - examples of 390
 - extended IP
 - configuring for QoS classification 429
 - creating 378
 - matching criteria 376
 - host keyword 378
 - IP
 - creating 376
 - implicit deny 378, 381, 383
 - implicit masks 378
 - management interfaces, applying to 387
 - matching criteria 372, 376
 - named 381
 - physical interfaces, applying to 388
 - undefined 387, 388
 - virtual terminal lines, setting on 388
 - MAC extended 386
 - matching 376
 - monitoring 389
 - named 381
 - numbers 376
 - protocol parameters 379
 - standard IP
 - configuring for QoS classification 428
 - creating 377
 - matching criteria 376
 - time ranges 383
 - unsupported features 376
- ACP
- system-defined mask 373
 - understanding 373
 - user-defined mask 373
- address resolution 116
- Address Resolution Protocol
- See ARP table
- addresses
- displaying the MAC address table 115
 - dynamic
 - accelerated aging 190
 - changing the aging time 111
 - default aging 190
 - defined 110
 - learning 111
 - removing 112
- MAC
- discovering 116
- multicast STP address management 190
- static
- adding and removing 114
 - defined 110
- advertisements
- CDP 331
 - VTP 258, 276, 277
- aggregated ports

- See EtherChannel
- aging time
 - accelerated
 - for MSTP 224
 - for STP 190, 202
 - MAC address table 111
 - maximum
 - for MSTP 224
 - for STP 203
- aging, accelerating 190
- alarms, RMON 354
- allowed-VLAN list 260
- ARP table
 - address resolution 116
 - managing 116
- attributes, RADIUS
 - vendor-proprietary 144
 - vendor-specific 143
- authentication
 - local mode with AAA 145
 - NTP associations 96
 - RADIUS
 - key 135
 - login 137
 - TACACS+
 - defined 125
 - key 127
 - login 128
- See also port-based authentication
- authoritative time source, described 94
- authorization
 - with RADIUS 141
 - with TACACS+ 126, 130
- authorized ports with 802.1X 152
- automatic discovery
 - adding member switches 86
 - considerations
 - beyond a non-candidate device 77, 78
 - brand new switches 78
 - connectivity 75
 - management VLANs 77, 78
 - non-CDP-capable devices 76
 - non-cluster-capable devices 76
 - creating a cluster standby group 87
 - in switch clusters 75
 - See also CDP
- automatic QoS
 - See QoS
- automatic recovery, clusters 79
 - See also HSRP
- autonegotiation
 - interface configuration guidelines 175
 - mismatches 457

B

- BackboneFast
 - described 231
 - enabling 239
 - support for 18
- bandwidth graphs 41

- banners
 - configuring
 - login 110
 - message-of-the-day login 109
 - default configuration 109
 - when displayed 108
- boot loader
 - described 63
 - environment variables 69
 - trap-door mechanism 63
- booting
 - boot loader, function of 63
 - boot process 63
 - manually 68
 - specific image 69
- BPDU
 - error-disabled state 228
 - filtering 228
 - RSTP format 213
- BPDU filtering
 - described 228
 - enabling 237
 - support for 19
- BPDU guard
 - described 228
 - enabling 236
 - support for 18
- broadcast storm control
 - configuring 315
 - disabling 316
- browser configuration 35, 73
- buttons, CMS 58

C

- cables, monitoring for unidirectional links 327
- candidate switch
 - adding 86
 - automatic discovery 75
 - defined 74
 - HC 88
 - passwords 86
 - requirements 74
 - standby group 87
 - See also command switch, cluster standby group, and member switch
- cautions 14
- CC (command switch) 88
- CDP
 - and trusted boundary 426
 - automatic discovery in switch clusters 75
 - configuring 331
 - default configuration 331
 - described 331
 - disabling for routing device 333
 - enabling and disabling
 - on a switch 333
 - on an interface 333
 - monitoring 334
 - overview 331
 - transmission timer and holdtime, setting 332

- updates 332
- CGMP, joining multicast group 290
- Cisco Discovery Protocol
 - See CDP
- CiscoWorks 2000 21, 398
- class maps for QoS
 - configuring 432
 - described 415
- clearing interfaces 180
- CLI
 - abbreviating commands 28
 - command modes 25
 - described 21
 - editing features
 - enabling and disabling 30
 - keystroke editing 31
 - wrapped lines 32
 - error messages 28
 - getting help 27
 - history
 - changing the buffer size 29
 - described 29
 - disabling 30
 - recalling commands 29
 - managing clusters 90
 - no and default forms of commands 28
- client mode, VTP 277
- clock
 - See system clock
- Cluster Management Suite
 - See CMS
- cluster standby group
 - automatic recovery 81
 - considerations 80
 - creating 87
 - defined 73
 - requirements 74
 - virtual IP address 80
 - See also HSRP
- cluster tree, described 39
- clusters, switch
 - accessing 82
 - adding member switches 86
 - automatic discovery 75
 - automatic recovery 79
 - command switch configuration 85
 - compatibility 75
 - creating 85
 - creating a cluster standby group 87
 - managing
 - through CLI 90
 - through SNMP 90
 - planning 75
 - planning considerations
 - automatic discovery 75
 - automatic recovery 79
 - CLI 90
 - host names 82
 - IP addresses 82
 - management VLAN 84
 - passwords 83
 - RADIUS 83
 - SNMP 83, 90
 - switch-specific features 85
 - TACACS+ 83
 - redundancy 87
 - verifying 89
 - See also candidate switch, command switch, cluster standby group, member switch, and standby command switch
- CMS
 - advantages 21
 - cluster tree 39
 - described 21, 35
 - displaying system messages 49
 - features 36
 - Front Panel images 40
 - Front Panel view 38
 - interaction modes 54
 - menu bar 45
 - online help 55
 - tool tips 55
 - toolbar 50
 - Topology view 42
 - window components 56
 - wizards 54
- command modes 25
- command switch
 - accessing 80
 - active (AC) 79, 88
 - command switch with HSRP disabled (CC) 88
 - defined 73
 - enabling 85
 - passive (PC) 79, 88
 - priority 79
 - recovery
 - from command-switch failure 79
 - redundant 79, 87
 - requirements 74
 - standby (SC) 79, 88
 - See also candidate switch, cluster standby group, member switch, and standby command switch
- command-line interface
 - See CLI
- commands
 - abbreviating 28
 - no and default 28
 - setting privilege levels 123
- community strings
 - configuring 83, 400
 - for cluster switches 398
 - in clusters 83
 - overview 397
 - SNMP 83
- config-vlan mode 26, 248
- config.text 68
- configuration files
 - clearing the startup configuration 484
 - creating using a text editor 475
 - default name 68
 - deleting a stored configuration 484
 - downloading
 - preparing 476, 478, 481

- reasons for 474
- using FTP 478
- using RCP 482
- using TFTP 476
- guidelines for creating and using 474
- invalid combinations when copying 470
- limiting TFTP server access 407
- specifying the filename 68
- system contact and location information 407
- types and location 475
- uploading
 - preparing 476, 478, 481
 - reasons for 474
 - using FTP 479
 - using RCP 483
 - using TFTP 477
- VMPS database 267
- configuration settings, saving 65
- configure terminal command 168
- connections, secure remote 147
- connectivity problems 457
- consistency checks in VTP version 2 278
- conventions
 - command 13
 - for examples 14
 - text 13
- CoS
 - configuring 416
 - configuring priority queues 438
 - defining 417
- CoS-to-DSCP map for QoS 436
- counters, clearing interface 180
- crashinfo file 463
- cryptographic software image 146

D

- daylight saving time 103
- debugging
 - enabling all system diagnostics 461
 - enabling for a specific feature 461
 - redirecting error message output 462
 - using commands 460
- default commands 28
- default configuration
 - 802.1X 156
 - banners 109
 - booting 67
 - CDP 331
 - DNS 107
 - EtherChannel 449
 - IGMP filtering 309
 - IGMP snooping 293
 - initial switch information 64
 - Layer 2 interfaces 173
 - MAC address table 111
 - MSTP 215
 - MVR 303
 - NTP 95
 - optional spanning-tree features 235
 - password and privilege level 118
- QoS 422
- RADIUS 134
- RMON 354
- RSPAN 342
- SNMP 399
- SPAN 342
- STP 193
- system message logging 361
- system name and prompt 106
- TACACS+ 127
- UDLD 328
- VLAN, Layer 2 Ethernet interfaces 257
- VLANs 249
- VMPS 269
- VTP 280
- default gateway 64
- deleting VLANs 251
- description command 177
- destination addresses, in ACLs 380
- detecting indirect link failures, STP 231
- device discovery protocol 331
- device labels 44
- Device Manager 38
 - See also Switch Manager
- device pop-up menu, Front Panel view 51
- Differentiated Services architecture, QoS 411
- Differentiated Services Code Point 412
- directories
 - changing 469
 - creating and removing 469
 - displaying the working 469
- discovery, clusters
 - See automatic discovery
- display options, Topology view 45
- Disqualification Code option 54
- DNS
 - default configuration 107
 - displaying the configuration 108
 - overview 107
 - setting up 107
- documentation
 - related 14
- Domain Name System
 - See DNS
- domain names
 - DNS 107
 - VTP 281
- downloading
 - configuration files
 - preparing 476, 478, 481
 - reasons for 474
 - using FTP 478
 - using RCP 482
 - using TFTP 476
 - image files
 - deleting old image 488
 - preparing 486, 489, 493
 - reasons for 484
 - using FTP 490
 - using RCP 494
 - using TFTP 487
- DSCP 412

- DSCP-to-CoS map for QoS 437
- DTP 19, 256
- duplex mode, configuring 174
- dynamic access mode 42
- dynamic access ports
 - characteristics 245
 - configuring 270
 - defined 166
- dynamic addresses
 - See addresses
- dynamic desirable trunking mode 256
- dynamic port VLAN membership
 - described 267
 - reconfirming 271
 - troubleshooting 273
 - types of connections 270
 - VMPS database configuration file 267
- Dynamic Trunking Protocol
 - See DTP

E

- editing features
 - enabling and disabling 30
 - keystrokes used 31
 - wrapped lines 32
- egress port scheduling 417
- enable password 119
- enable secret password 119
- encapsulation 416
- encryption for passwords 119
- environment variables
 - function of 70
 - location in Flash 69
- error messages
 - during command entry 28
 - setting the display destination device 362
 - severity levels 365
 - system message format 360
- EtherChannel
 - automatic creation of 444
 - configuration guidelines 449
 - default configuration 449
 - destination MAC address forwarding 447
 - displaying status 455
 - forwarding methods 452
 - interaction with STP 449
 - Layer 2 interfaces, configuring 450
 - load balancing 446, 452
 - number of interfaces per 443
 - overview 443
 - PAgP
 - aggregate-port learners 446
 - displaying status 455
 - interaction with other features 446
 - learn method and priority configuration 453
 - modes 444
 - overview 444
 - silent mode 445
 - support for 17
 - port groups 167

- port-channel interfaces
 - described 444
 - numbering of 444
 - source MAC address forwarding 447
- EtherChannel guard
 - described 233
 - enabling 239
- Ethernet VLANs
 - adding 249
 - defaults and ranges 249
 - modifying 249
- events, RMON 354
- examples
 - conventions for 14
- expedite queue, QoS 417
- expert mode 54
- extended system ID
 - MSTP 218
 - STP 185, 196
- extended-range VLANs
 - configuration guidelines 253
 - configuring 252
 - creating 253
 - defined 243
- Extensible Authentication Protocol over LAN 149

F

- fallback VLAN name 267
- features, IOS 17
- file system
 - displaying available file systems 467
 - displaying file information 469
 - local file system names 467
 - network file system names 470
 - setting the default 468
- files
 - copying 470
 - deleting 471
 - displaying the contents of 473
 - tar
 - creating 471
 - displaying the contents of 472
 - extracting 473
 - image file format 485
- files, crashinfo
 - description 463
 - displaying the contents of 463
 - location 463
- filtering show and more command output 33
- filters, IP
 - See ACLs, IP
- Flash device, number of 467
- flow control 176
- forward-delay time
 - MSTP 224
 - STP 187, 202
- forwarding
 - See broadcast storm control
- Front Panel images, CMS 40
- Front Panel view

- cluster tree 39
- described 38
- pop-up menus 51
- port icons 41
- port LEDs 41
- switch images 40

FTP

- configuration files
 - downloading 478
 - overview 477
 - preparing the server 478
 - uploading 479
- image files
 - deleting old image 492
 - downloading 490
 - preparing the server 489
 - uploading 492

G

- get-bulk-request operation 397
- get-next-request operation 397, 398
- get-request operation 397, 398
- get-response operation 397
- global configuration mode 26
- graphs, bandwidth 41
- guide
 - audience 13
 - purpose 13
- guide mode 54

H

- HC (candidate switch) 88
- hello time
 - MSTP 223
 - STP 202
- Help Contents 55
- help, for the command line 27
- history
 - changing the buffer size 29
 - described 29
 - disabling 30
 - recalling commands 29
- history table, level and number of syslog messages 367
- host name list, CMS 57
- host names
 - abbreviations appended to 88
 - in clusters 82
- hosts, limit on dynamic ports 273
- HP OpenView 21
- HSRP
 - automatic cluster recovery 81
 - cluster standby group considerations 80
 - See also clusters, cluster standby group, and standby command switch

I

- ICMP ping
 - executing 458
 - overview 458
- icons
 - colors
 - cluster tree 40
 - Topology view 44
 - Front Panel view 41
 - toolbar 50
 - Topology view 43
- IDS, using with SPAN and RSPAN 338
- IGMP
 - join messages 290
 - joining multicast group 290
 - leave processing, enabling 297
 - leaving multicast group 292
 - queries 290
- IGMP filtering
 - configuring 309
 - default configuration 309
 - described 308
 - monitoring 312
- IGMP groups, setting the maximum number 311
- IGMP profile
 - applying 310
 - configuration mode 309
 - configuring 309
- IGMP snooping
 - configuring 293
 - default configuration 293
 - definition 289
 - enabling and disabling 293
 - global configuration 293
 - Immediate Leave 292
 - method 294
 - monitoring 299
 - VLAN configuration 294
- Immediate-Leave, IGMP 292
- ingress port scheduling 417
- interaction modes, CMS 54
- interface
 - number 168
 - range macros 171
- interface command 168
- interface configuration mode 26
- interfaces
 - configuration guidelines 175
 - configuring 168
 - configuring duplex mode 174
 - configuring speed 174
 - counters, clearing 180
 - described 177
 - descriptive name, adding 177
 - displaying information about 178
 - flow control 176
 - IOS supported 21
 - monitoring 178
 - naming 177
 - physical, identifying 168

- range of 169
- restarting 181
- shutting down 181
- supported 173
- types of 165
- interfaces range macro command 171
- Intrusion Detection System
 - See IDS
- inventory, cluster 89
- IOS command-line interface
 - See CLI
- IP
 - named extended ACL 382
 - named standard ACL 381
 - numbered extended ACL 378
 - numbered standard ACL 377
- IP ACLs
 - applying to
 - management interfaces 387
 - physical interfaces 388
 - extended, creating 378
 - for QoS classification 428
 - implicit deny 378, 381, 383
 - implicit masks 378
 - management interfaces, applying to 387
 - named 381
 - physical interfaces, applying to 388
 - standard, creating 377
 - undefined 387, 388
 - virtual terminal lines, setting on 388
- IP addresses
 - candidate or member 74, 82
 - cluster access 74
 - command switch 74, 80, 82
 - discovering 116
 - management VLAN 84
 - redundant clusters 80
 - standby command switch 80, 82
 - See also IP information
- ip igmp profile command 309
- IP information
 - assigned
 - manually 64
 - default configuration 64
- IP multicast routing and IGMP snooping 289, 293
- IP phones
 - automatic classification and queueing 418
 - trusted boundary for QoS 426
- IP protocols in ACLs 380

J

- Java plug-in configuration 35, 73
- join messages, IGMP 290

L

- LACP
 - See EtherChannel
- Layer 2 frames, classification with CoS 412

- Layer 2 interfaces, default configuration 173
- Layer 2 traceroute
 - and ARP 460
 - and CDP 459
 - described 459
 - IP addresses and subnets 460
 - MAC addresses and VLANs 460
 - multicast traffic 460
 - multiple devices on a port 460
 - supported switches 459
 - unicast traffic 459
 - usage guidelines 459
- Layer 2 trunks 256
- Layer 3 packets, classification methods 412
- Layer 3 parameters of ACEs 379
- Layer 4 parameters of ACEs 379
- leave processing, IGMP 297
- LEDs
 - port 41
- legend, CMS icons and labels 50
- line configuration mode 26
- link labels 44
- link pop-up menu, Topology view 52
- links, unidirectional 327
- lists, CMS 57
- log messages
 - See system message logging
- login authentication
 - with RADIUS 137
 - with TACACS+ 128
- login banners 108
- loop guard
 - described 234
 - enabling 240
 - support for 19

M

- MAC address multicast entries, monitoring 299
- MAC address-to-VLAN mapping 266
- MAC addresses
 - adding
 - sticky secure 319
 - aging time 111
 - and VLAN association 111
 - building the address table 111
 - default configuration 111
 - discovering 116
 - displaying 115
 - dynamic
 - learning 111
 - removing 112
 - in ACLs 386
 - static
 - adding 115
 - characteristics of 114
 - removing 115
- MAC extended access lists 386
- management options
 - benefits
 - clustering 21

- CMS 21
- CLI 25
- CMS 35
- overview 21
- management VLAN
 - changing 84
 - considerations in switch clusters 77, 78, 84
 - discovery through different management VLANs 78
 - discovery through same management VLAN 77
 - IP address 84
- mapping tables for QoS
 - configuring
 - DSCP 436
 - DSCP-to-CoS 437
 - described 414
- matching, ACLs 376
- maximum aging time
 - MSTP 224
 - STP 203
- maximum hop count, MSTP 224
- member switch
 - adding 86
 - automatic discovery 75
 - defined 73
 - managing 90
 - passwords 82
 - requirements 74
 - See also candidate switch, cluster standby group, and standby command switch
- membership mode, VLAN port 245
- menu bar
 - described 45
 - variations 45
- messages
 - system 49
 - to users through banners 108
- MIBs
 - overview 395
 - SNMP interaction with 398
 - supported 465
- mirroring traffic for analysis 337
- mismatches, autonegotiation 457
- modes, port 41
- monitoring
 - access groups 389
 - ACLs 389
 - cables for unidirectional links 327
 - CDP 334
 - IGMP
 - filters 312
 - snooping 299
 - interfaces 178
 - multicast router interfaces 299
 - MVR 307
 - network traffic for analysis with probe 337
 - port protection 325
 - speed and duplex mode 175
 - traffic flowing among switches 353
 - traffic suppression 325
 - VLANs 254
 - VMPS 272

- VTP 288
- MSTP
 - boundary ports
 - configuration guidelines 216
 - described 209
 - BPDU filtering
 - described 228
 - enabling 237
 - BPDU guard
 - described 228
 - enabling 236
 - CIST, described 206
 - configuration guidelines 216, 235
 - configuring
 - forward-delay time 224
 - hello time 223
 - link type for rapid convergence 225
 - maximum aging time 224
 - maximum hop count 224
 - MST region 217
 - path cost 221
 - port priority 220
 - root switch 218
 - secondary root switch 219
 - switch priority 222
- CST
 - defined 206
 - operations between regions 207
- default configuration 215
- default optional feature configuration 235
- displaying status 226
- enabling the mode 217
- EtherChannel guard
 - described 233
 - enabling 239
- extended system ID
 - effects on root switch 218
 - effects on secondary root switch 219
 - unexpected behavior 218
- instances supported 191
- interface state, blocking to forwarding 227
- interoperability and compatibility among modes 191
- interoperability with 802.1D
 - described 209
 - restarting migration process 225
- IST
 - defined 206
 - master 207
 - operations within a region 207
- loop guard
 - described 234
 - enabling 240
- mapping VLANs to MST instance 217
- MST region
 - CIST 206
 - configuring 217
 - described 206
 - hop-count mechanism 208
 - IST 206
 - supported spanning-tree instances 206
- overview 205

- Port Fast
 - described 227
 - enabling 235
- preventing root switch selection 233
- root guard
 - described 233
 - enabling 240
- root switch
 - configuring 218
 - effects of extended system ID 218
 - unexpected behavior 218
- shutdown Port Fast-enabled port 228
- multicast groups
 - and IGMP snooping 293
 - Immediate Leave 292
 - joining 290
 - leaving 292
 - static joins 296
- multicast router interfaces, monitoring 299
- multicast router ports, adding 295
- Multicast VLAN Registration
 - See MVR
- Multiple Spanning Tree Protocol
 - See MSTP
- MVR
 - configuring interfaces 305
 - default configuration 303
 - described 301
 - modes 305
 - monitoring 307
 - setting global parameters 304

N

- named IP ACLs 381
- native VLAN
 - configuring 262
 - default 262
- negotiate trunk mode 42
- network management
 - CDP 331
 - RMON 353
 - SNMP 395
- Network Time Protocol
 - See NTP
- no commands 28
- nontrunking mode 256
- normal-range VLANs
 - configuration modes 247
 - defined 243
- NTP
 - associations
 - authenticating 96
 - defined 94
 - enabling broadcast messages 98
 - peer 97
 - server 97
 - default configuration 95
 - displaying the configuration 102
 - overview 94
 - restricting access

- creating an access group 99
 - disabling NTP services per interface 101
- source IP address, configuring 101
- stratum 94
- synchronizing devices 97
- time
 - services 94
 - synchronizing 94

O

- online help 55
- out-of-profile markdown 20

P

- PAGP
 - See EtherChannel
- pass-through mode 428
- passwords
 - default configuration 118
 - encrypting 119
 - in clusters 83, 86
 - overview 117
 - setting
 - enable 118
 - enable secret 119
 - Telnet 121
 - with usernames 121
 - VTP domain 281
- path cost
 - MSTP 221
 - STP 199
- PC (passive command switch) 79, 88
- per-VLAN spanning-tree plus
 - See PVST+
- physical ports 165
- PIM-DVMRP, as snooping method 294
- ping
 - character output description 458
 - executing 458
 - overview 458
- policers
 - configuring for each matched traffic class 433
 - described 413
 - number of 20, 416
 - types of 415
- policing 20, 413
- policy maps for QoS
 - characteristics of 433
 - configuring 433
 - described 415
 - displaying 440
- Port Aggregation Protocol
 - See EtherChannel
 - See PAGP
- Port Fast
 - described 227
 - enabling 235
 - mode, spanning tree 269

- support for 18
- port icons, Front Panel view 41
- port LEDs
 - port modes 41
- port membership modes, VLAN 244
- port modes, described 41
- port pop-up menu, Front Panel view 51
- port priority
 - MSTP 220
 - STP 198
- port scheduling 417
- port security
 - aging 323
 - configuring 321
 - default configuration 320
 - described 318
 - displaying 325
 - sticky learning 319
 - violations 319
 - with other features 320
- port-based authentication
 - authentication server
 - defined 150
 - RADIUS server 150
 - client, defined 150
 - configuration guidelines 156
 - configuring
 - 802.1X authentication 157
 - guest VLAN 163
 - host mode 162
 - manual re-authentication of a client 160
 - periodic re-authentication 160
 - quiet period 161
 - RADIUS server 160
 - RADIUS server parameters on the switch 159
 - switch-to-client frame-retransmission number 162
 - switch-to-client retransmission time 161
 - default configuration 156
 - described 149
 - device roles 149
 - displaying statistics 164
 - EAP-request/identity frame 151
 - EAP-response/identity frame 151
 - EAPOL-start frame 151
 - enabling
 - 802.1X with guest VLAN 155
 - 802.1X with port security 153, 163
 - 802.1X with VLAN assignment 154, 157
 - encapsulation 150
 - initiation and message exchange 151
 - method lists 157
 - ports
 - authorization state and dot1x port-control command 152
 - authorized and unauthorized 152
 - resetting to default values 164
 - switch
 - as proxy 150
 - RADIUS client 150
 - topologies, supported 152

- port-channel
 - See EtherChannel
- port-shutdown response, VMPS 266
- ports
 - 802.1Q trunk 42
 - access 165
 - dynamic access 42, 245
 - negotiate trunk 42
 - priority 416
 - protected 317
 - secure 318
 - static-access 42, 245, 251
 - switch 27, 165
 - trunks 255
 - VLAN assignments 251
- preferential treatment of traffic
 - See QoS
- preventing unauthorized access 117
- priority
 - port, described 416
- private VLAN edge ports
 - See protected ports
- privilege levels
 - changing the default for lines 123
 - exiting 124
 - logging into 124
 - overview 117, 122
 - setting a command with 123
- privileged EXEC mode 26
- protected ports 18, 317
- pruning, VTP
 - enabling 286
 - enabling on a port 261
 - examples 279
 - overview 278
- pruning-eligible list
 - changing 261
 - for VTP pruning 278
 - VLANs 287
- publications, related 14
- PVRST+ 244
- PVST+ 244
 - 802.1Q trunking interoperability 192
 - described 190
 - instances supported 191

Q

- QoS
 - auto-QoS
 - configuration and defaults display 421
 - displaying 421
 - effects on NVRAM configuration 420
 - basic model 413
 - class maps
 - configuring 432
 - classification
 - class maps, described 415
 - defined 413
 - in frames and packets 412
 - IP ACLs, described 414

- MAC ACLs, described 414
- pass-through mode, described 428
- policy maps, described 415
- port default, described 413
- trust DSCP, described 414
- trusted boundary, described 426
- trusted CoS, described 414
- types for IP traffic 414
- types for non-IP traffic 413
- configuration examples
 - common wiring closet 441
 - intelligent wiring closet 442
- configuration guidelines 422
- configuring
 - class maps 432
 - CoS and WRR 438
 - default port CoS value 425
 - egress queues 438
 - IP extended ACLs 429
 - IP standard ACLs 428
 - MAC ACLs 431
 - policy maps 433
 - port trust states within the domain 423
 - QoS policy 428
 - trusted boundary 426
- default configuration 422
- displaying statistics 440
- egress port scheduling 417
- enabling expedite queue 439
- expedite queue
 - described 417
 - enabling 439
- ingress port scheduling 417
- IP phones
 - automatic classification and queueing 418
 - detection and trusted settings 418
- IP phones, detection and trusted settings 426
- mapping tables
 - CoS-to-DSCP 436
 - displaying 440
 - DSCP-to-CoS 437
 - types of 414
- marked-down actions 435
- marking, described 413, 415
- overview 411
- pass-through mode 428
- policers
 - configuring 435
 - described 415
 - number of 416
 - types of 415
- policing, described 413, 415
- policy maps
 - characteristics of 433
 - configuring 433
 - displaying 440
- queueing, defined 413
- scheduling, defined 413
- support for 20
- trust states 413
- trusted boundary 426
- understanding 411

- quality of service
 - See QoS
- queries, IGMP 290

R

- RADIUS
 - attributes
 - vendor-proprietary 144
 - vendor-specific 143
 - configuring
 - accounting 142
 - authentication 137
 - authorization 141
 - communication, global 135, 142
 - communication, per-server 134, 135
 - multiple UDP ports 134
 - default configuration 134
 - defining AAA server groups 139
 - displaying the configuration 145
 - identifying the server 134
 - in clusters 83
 - limiting the services to the user 141
 - method list, defined 133
 - operation of 133
 - overview 132
 - suggested network environments 132
 - tracking services accessed by user 142
- range
 - macro 171
 - of interfaces 170
- rapid convergence 211
- rapid per-VLAN spanning-tree plus
 - See rapid PVST+
- rapid PVST+
 - 802.1Q trunking interoperability 192
 - described 190
 - instances supported 191
- Rapid Spanning Tree Protocol
 - See RSTP
- rapid-PVST+ 244
- rcommand command 90
- RCP
 - configuration files
 - downloading 482
 - overview 480
 - preparing the server 481
 - uploading 483
 - image files
 - deleting old image 496
 - downloading 494
 - preparing the server 493
 - uploading 496
- reconfirmation interval, VMPS, changing 271
- redundancy
 - EtherChannel 444
 - STP
 - backbone 189
 - path cost 264
 - port priority 263
- redundant clusters

- See cluster standby group
- redundant links and UplinkFast 238
- reloading software 71
- Remote Authentication Dial-In User Service
 - See RADIUS
- Remote Copy Protocol
 - See RCP
- remote monitoring
 - see RMON
- Remote Network Monitoring
 - See RMON
- resetting a UDLD-shutdown interface 330
- restricting access
 - NTP services 99
 - overview 117
 - passwords and privilege levels 117
 - RADIUS 131
 - TACACS+ 124
- retry count, VMPS, changing 272
- RFC
 - 1112, IP multicast and IGMP 289
 - 1157, SNMPv1 395
 - 1305, NTP 94
 - 1757, RMON 353
 - 1901, SNMPv2C 396
 - 1902 to 1907, SNMPv2 396
 - 2236, IP multicast and IGMP 289
 - 2273-2275, SNMPv3 396
- RMON
 - default configuration 354
 - displaying status 357
 - enabling alarms and events 354
 - groups supported 353
 - overview 353
 - statistics
 - collecting group Ethernet 356
 - collecting group history 356
- root guard
 - described 233
 - enabling 240
 - support for 19
- root switch
 - MSTP 218
 - STP 196
- RSPAN
 - configuration guidelines 347
 - default configuration 342
 - destination ports 339
 - displaying status 351
 - IDS 338
 - interaction with other features 341
 - monitored ports 339
 - monitoring ports 339
 - overview 337
 - received traffic 338
 - reflector port 340
 - session limits 342
 - sessions
 - creating 348
 - defined 338
 - removing source (monitored) ports 350
 - specifying monitored ports 348

- source ports 339
- transmitted traffic 339
- RSTP
 - active topology, determining 210
 - BPDU
 - format 213
 - processing 214
 - designated port, defined 210
 - designated switch, defined 210
 - interoperability with 802.1D
 - described 209
 - restarting migration process 225
 - topology changes 214
 - overview 210
 - port roles
 - described 210
 - synchronized 212
 - proposal-agreement handshake process 211
 - rapid convergence
 - described 211
 - edge ports and Port Fast 211
 - point-to-point links 211, 225
 - root ports 211
 - root port, defined 210
 - See also MSTP
- running configuration, saving 65

S

- SC (standby command switch) 79, 88
- scheduled reloads 71
- secure ports, configuring 318
- secure remote connections 147
- Secure Shell
 - See SSH
- security, port 318
- sequence numbers in log messages 365
- server mode, VTP 276
- service-provider network, MSTP and RSTP 205
- set-request operation 398
- severity levels, defining in system messages 365
- show and more command output, filtering 33
- show cdp traffic command 334
- show cluster members command 90
- show configuration command 177
- show interfaces command 175, 177
- show running-config command
 - displaying ACLs 387, 388
 - interface description in 177
- shutdown command on interfaces 181
- Simple Network Management Protocol
 - See SNMP
- SNAP 331
- SNMP
 - accessing MIB variables with 398
 - agent
 - described 397
 - disabling 400
 - community strings
 - configuring 400
 - for cluster switches 398

- overview 397
 - configuration examples 408
 - default configuration 399
 - groups 402
 - in clusters 83
 - informs
 - and trap keyword 404
 - described 398
 - differences from traps 398
 - enabling 406
 - limiting access by TFTP servers 407
 - limiting system log messages to NMS 366
 - manager functions 397
 - managing clusters with 90
 - MIBs
 - supported 465
 - notifications 398
 - overview 395, 398
 - status, displaying 409
 - system contact and location 407
 - trap manager, configuring 405
 - traps
 - described 397, 398
 - differences from informs 398
 - enabling 404
 - enabling MAC address notification 112
 - overview 395, 398
 - types of 404
 - users 402
 - versions supported 395
- snooping, IGMP 289
- software images
 - location in Flash 485
 - scheduling reloads 71
 - tar file format, described 485
- source addresses, in ACLs 380
- SPAN
 - configuration guidelines 342
 - default configuration 342
 - destination ports 339
 - displaying status 351
 - IDS 338
 - interaction with other features 341
 - monitored ports 339
 - monitoring ports 339
 - overview 20, 337
 - received traffic 338
 - session limits 342
 - sessions
 - creating 343
 - defined 338
 - removing destination (monitoring) ports 346
 - removing source (monitored) ports 346
 - specifying monitored ports 343
 - source ports 339
 - transmitted traffic 339
- spanning tree and native VLANs 257
- Spanning Tree Protocol
 - See STP
- speed, configuring on interfaces 174
- SSH
 - configuring 147
 - cryptographic software image 146
 - described 147
 - displaying settings 147
- Standby Command Configuration window 89
- standby command switch
 - configuring 87
 - considerations 80
 - defined 73
 - priority 79
 - requirements 74
 - virtual IP address 80
 - See also cluster standby group and HSRP
- standby group, cluster
 - See cluster standby group and HSRP
- startup configuration
 - booting
 - manually 68
 - specific image 69
 - clearing 484
 - configuration file
 - specifying the filename 68
 - default boot configuration 67
- static access mode 42
- static access ports
 - assigning to VLAN 251
 - defined 166, 245
- static addresses
 - See addresses
- static VLAN membership 244
- statistics
 - 802.1X 164
 - CDP 334
 - interface 179
 - QoS ingress and egress 440
 - RMON group Ethernet 356
 - RMON group history 356
 - SNMP input and output 409
 - VTP 288
- sticky learning
 - configuration file 319
 - defined 319
 - disabling 319
 - enabling 319
 - saving addresses 319
- storm control
 - described 315
 - displaying 325
- STP
 - accelerating root port selection 229
 - BackboneFast
 - described 231
 - enabling 239
 - BPDU filtering
 - described 228
 - enabling 237
 - BPDU guard
 - described 228
 - enabling 236
 - BPDU message exchange 184
 - configuration guidelines 193, 235
 - configuring
 - forward-delay time 202

- hello time 202
- maximum aging time 203
- path cost 199
- port priority 198
- root switch 196
- secondary root switch 197
- spanning-tree mode 194
- switch priority 201
- counters, clearing 203
- default configuration 193
- default optional feature configuration 235
- designated port, defined 185
- designated switch, defined 185
- detecting indirect link failures 231
- disabling 195
- displaying status 203
- EtherChannel guard
 - described 233
 - enabling 239
- extended system ID
 - affects on root switch 196
 - affects on the secondary root switch 197
 - overview 185
 - unexpected behavior 196
- features supported 18
- inferior BPDU 185
- instances supported 191
- interface state, blocking to forwarding 227
- interface states
 - blocking 187
 - disabled 188
 - forwarding 187, 188
 - learning 188
 - listening 188
 - overview 186
- interoperability and compatibility among modes 191
- limitations with 802.1Q trunks 191
- load sharing
 - overview 262
 - using path costs 264
 - using port priorities 263
- loop guard
 - described 234
 - enabling 240
- modes supported 190
- multicast addresses, affect of 190
- overview 183
- path costs 264, 265
- Port Fast
 - described 227
 - enabling 235
- port priorities 263
- preventing root switch selection 233
- protocols supported 190
- redundant connectivity 189
- root guard
 - described 233
 - enabling 240
- root port, defined 185
- root switch
 - affects of extended system ID 185, 196
 - configuring 196
 - election 185
 - unexpected behavior 196
- shutdown Port Fast-enabled port 228
- superior BPDU 185
- timers, described 201
- UplinkFast
 - described 229
 - enabling 238
- stratum, NTP 94
- summer time 103
- SunNet Manager 21
- Switch Manager 38
 - See also Device Manager
- switch ports 27, 165
- switch priority
 - MSTP 222
 - STP 201
- switchport protected command 317
- syslog
 - See system message logging
- system clock
 - configuring
 - daylight saving time 103
 - manually 102
 - summer time 103
 - time zones 103
 - displaying the time and date 103
 - overview 93
 - See also NTP
- system message logging
 - default configuration 361
 - defining error message severity levels 365
 - disabling 361
 - displaying the configuration 369
 - enabling 361
 - facility keywords, described 368
 - level keywords, described 366
 - limiting messages 366
 - message format 360
 - overview 359
 - sequence numbers, enabling and disabling 365
 - setting the display destination device 362
 - synchronizing log messages 363
 - timestamps, enabling and disabling 364
 - UNIX syslog servers
 - configuring the daemon 367
 - configuring the logging facility 368
 - facilities supported 368
- system messages on CMS 49
- system name
 - default configuration 106
 - default setting 106
 - manual configuration 106
 - See also DNS
- system prompt
 - default setting 105, 106
 - manual configuration 106

T

- tables, CMS 57
- tabs, CMS 57
- TACACS+
 - accounting, defined 126
 - authentication, defined 125
 - authorization, defined 126
 - configuring
 - accounting 131
 - authentication key 127
 - authorization 130
 - login authentication 128
 - default configuration 127
 - displaying the configuration 131
 - identifying the server 127
 - in clusters 83
 - limiting the services to the user 130
 - operation of 126
 - overview 125
 - tracking services accessed by user 131
- tar files
 - creating 471
 - displaying the contents of 472
 - extracting 473
 - image file format 485
- Telnet
 - accessing management interfaces 33
 - accessing the CLI 21
 - from a browser 33
 - setting a password 121
- Terminal Access Controller Access Control System Plus
 - See TACACS+
- terminal lines, setting a password 121
- TFTP
 - configuration files
 - downloading 476
 - preparing the server 476
 - uploading 477
 - image files
 - deleting 488
 - downloading 487
 - preparing the server 486
 - uploading 488
 - limiting access by servers 407
- time
 - See NTP and system clock
- time ranges in ACLs 383
- time zones 103
- time-range command 383
- timestamps in log messages 364
- Token Ring VLANs
 - support for 246
 - VTP support 278
- tool tips 55
- toolbar 50
- Topology view
 - described 36, 42
 - device icons 44
 - device labels 44
 - display options 45
 - icons 43
 - link labels 44
- traceroute, Layer 2
 - and ARP 460
 - and CDP 459
 - described 459
 - IP addresses and subnets 460
 - MAC addresses and VLANs 460
 - multicast traffic 460
 - multiple devices on a port 460
 - supported switches 459
 - unicast traffic 459
 - usage guidelines 459
- trademarks 501
- traffic
 - fragmented 372
 - unfragmented 372
- traffic policing 20
- transparent mode, VTP 277, 285
- trap-door mechanism 63
- traps
 - configuring MAC address notification 112
 - configuring managers 404
 - defined 397
 - enabling 112, 404
 - notification types 404
 - overview 395, 398
- troubleshooting
 - connectivity problems 457
 - detecting unidirectional links 327
 - displaying crash information 463
 - with CiscoWorks 398
 - with debug commands 460
 - with ping 458
 - with system message logging 359
- trunk ports
 - configuring 259
 - defined 166
- trunks
 - allowed-VLAN list 260
 - load sharing
 - setting STP path costs 264
 - using STP port priorities 263
 - native VLAN for untagged traffic 262
 - parallel 264
 - pruning-eligible list 261
 - to non-DTP device 256
 - understanding 256
 - VLAN 1 minimization 260
- trusted boundary 426
- twisted-pair Ethernet, detecting unidirectional links 327

U

- UDLD
 - default configuration 328
 - echoing detection mechanism 328
 - enabling
 - globally 329

- per interface 329
 - link-detection mechanism 327
 - neighbor database 327
 - overview 327
 - resetting an interface 330
 - status, displaying 330
- unauthorized ports with 802.1X 152
- UniDirectional Link Detection protocol
 - See UDLD
- UNIX syslog servers
 - daemon configuration 367
 - facilities supported 368
 - message logging configuration 368
- unrecognized Type-Length-Value (TLV) support 278
- upgrading software images
 - See downloading
- UplinkFast
 - described 229
 - enabling 238
 - support for 18
- uploading
 - configuration files
 - preparing 476, 478, 481
 - reasons for 474
 - using FTP 479
 - using RCP 483
 - using TFTP 477
 - image files
 - preparing 486, 489, 493
 - reasons for 484
 - using FTP 492
 - using RCP 496
 - using TFTP 488
- user EXEC mode 26
- username-based authentication 121

V

- version-dependent transparent mode 278
- virtual IP address
 - cluster standby group 80, 88
 - command switch 80, 88
 - See also IP addresses
- VLAN 1 minimization, support for 19
- VLAN configuration
 - at bootup 248
 - saving 248
- VLAN configuration mode 26, 248
- VLAN database
 - and startup configuration file 248
 - and VTP 275
 - VLAN configuration saved in 248
 - VLANs saved in 245
- vlan database command 248
- vlan global configuration command 248
- VLAN ID, discovering 116
- VLAN management domain 275
- VLAN Management Policy Server
 - See VMPS
- VLAN membership
 - confirming 271

- modes 245
- VLAN Query Protocol
 - See VQP
- VLAN Trunking Protocol
 - See VTP
- VLAN trunks 255, 256
- vlan.dat file 246
- VLANs
 - adding 249
 - adding to VLAN database 249
 - aging dynamic addresses 190
 - allowed on trunk 260
 - and spanning-tree instances 244, 247, 253
 - configuration guidelines, normal-range VLANs 247
 - configuration options 247
 - configuring 243
 - configuring IDs 1006 to 4094 253
 - creating in config-vlan mode 249
 - creating in VLAN configuration mode 250
 - default configuration 249
 - deleting 251
 - described 166, 243
 - displaying 254
 - extended-range 243, 252
 - illustrated 244
 - modifying 249
 - native, configuring 262
 - normal-range 243, 245
 - parameters 246
 - port membership modes 244
 - static-access ports 251
 - STP and 802.1Q trunks 191
 - supported 244
 - Token Ring 246
 - trunks, VLAN 1 minimization 260
 - VTP modes 276
- VMPS
 - administering 272
 - configuration example 273
 - configuration guidelines 269
 - default configuration 269
 - description 266
 - dynamic port membership
 - described 267
 - reconfirming 271
 - troubleshooting 273
 - entering server address 270
 - mapping MAC addresses to VLANs 266
 - monitoring 272
 - reconfirmation interval, changing 271
 - reconfirming membership 271
 - retry count, changing 272
- VQP 266
- VTP
 - adding a client to a domain 287
 - advertisements 258, 277
 - and extended-range VLANs 275
 - and normal-range VLANs 275
 - client mode, configuring 284
 - configuration
 - global configuration mode 280

- guidelines 281
 - privileged EXEC mode 280
 - requirements 282
 - saving 281
 - VLAN configuration mode 281
- configuration mode options 280
- configuration requirements 282
- configuration revision number
 - guideline 287
 - resetting 288
- configuring
 - client mode 284
 - server mode 282
 - transparent mode 285
- consistency checks 278
- default configuration 280
- described 275
- disabling 285
- domain names 281
- domains 275
- modes
 - client 277, 284
 - server 276, 282
 - transitions 276
 - transparent 277, 285
- monitoring 288
- passwords 281
- pruning
 - disabling 287
 - enabling 286
 - examples 279
 - overview 278
- pruning-eligible list, changing 261
- server mode, configuring 282
- statistics 288
- Token Ring support 278
- transparent mode, configuring 285
- using 275
- version 1 278
- version 2
 - configuration guidelines 282
 - disabling 286
 - enabling 286
 - overview 278
- version, guidelines 282

W

- warnings 14
- web-based management software
 - See CMS
- Weighted Round Robin
 - See WRR
- window components, CMS 56
- wizards 54
- WRR
 - configuring 439
 - defining 417
 - description 417



Part Number: 25K8411