



Catalyst 3750 Metro Switch Software Configuration Guide

Cisco IOS Release 12.1(14)AX
January 2004

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Customer Order Number: DOC-7815870=
Text Part Number: 78-15870-01



THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCIP, CCSP, the Cisco Arrow logo, the Cisco *Powered* Network mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0304R)

Catalyst 3750 Metro Switch Software Configuration Guide
Copyright © 2004 Cisco Systems, Inc. All rights reserved.



Preface xxxiii

Audience xxxiii

Purpose xxxiii

Conventions xxxiv

Related Publications xxxv

Obtaining Documentation xxxv

Cisco.com xxxv

Ordering Documentation xxxvi

Documentation Feedback xxxvi

Obtaining Technical Assistance xxxvi

Cisco TAC Website xxxvi

Opening a TAC Case xxxvii

TAC Case Priority Definitions xxxvii

Obtaining Additional Publications and Information xxxviii

CHAPTER 1

Overview 1-1

Features 1-1

Default Settings After Initial Switch Configuration 1-8

Network Configuration Examples 1-11

Multidwelling or Ethernet-to-the Subscriber Network 1-11

Ethernet Broadband Aggregation Network 1-13

Layer 2 VPN Application 1-14

Layer 3 VPN Application 1-15

Where to Go Next 1-16

CHAPTER 2

Using the Command-Line Interface 2-1

Understanding Command Modes 2-1

Understanding the Help System 2-3

Understanding Abbreviated Commands 2-4

Understanding no and default Forms of Commands 2-4

Understanding CLI Error Messages 2-5

Using Command History	2-5
Changing the Command History Buffer Size	2-5
Recalling Commands	2-6
Disabling the Command History Feature	2-6
Using Editing Features	2-6
Enabling and Disabling Editing Features	2-7
Editing Commands through Keystrokes	2-7
Editing Command Lines that Wrap	2-8
Searching and Filtering Output of show and more Commands	2-9
Accessing the CLI	2-9
Accessing the CLI through a Console Connection or through Telnet	2-10
Accessing the CLI from a Browser	2-10

CHAPTER 3

Assigning the Switch IP Address and Default Gateway	3-1
Understanding the Boot Process	3-1
Assigning Switch Information	3-2
Default Switch Information	3-3
Understanding DHCP-Based Autoconfiguration	3-3
DHCP Client Request Process	3-4
Configuring DHCP-Based Autoconfiguration	3-4
Configuring the DHCP Server	3-5
Configuring the TFTP Server	3-5
Configuring the DNS	3-6
Configuring the Relay Device	3-6
Obtaining Configuration Files	3-7
Example Configuration	3-8
Manually Assigning IP Information	3-9
Checking and Saving the Running Configuration	3-10
Modifying the Startup Configuration	3-10
Default Boot Configuration	3-11
Automatically Downloading a Configuration File	3-11
Specifying the Filename to Read and Write the System Configuration	3-11
Booting Manually	3-12
Booting a Specific Software Image	3-12
Controlling Environment Variables	3-13
Scheduling a Reload of the Software Image	3-15
Configuring a Scheduled Reload	3-15
Displaying Scheduled Reload Information	3-16

CHAPTER 4

Configuring IE2100 CNS Agents	4-1
Understanding IE2100 Series Configuration Registrar Software	4-1
CNS Configuration Service	4-2
CNS Event Service	4-3
NameSpace Mapper	4-3
Configuration Registrar and ConfigID, DeviceID, and Host Name	4-3
ConfigID	4-3
DeviceID	4-4
Host Name and DeviceID	4-4
Using Host Name, DeviceID, and ConfigID	4-4
Understanding CNS Embedded Agents	4-5
Initial Configuration	4-5
Incremental (Partial) Configuration	4-6
Synchronized Configuration	4-6
Configuring CNS Embedded Agents	4-6
Enabling Automated CNS Configuration	4-6
Enabling the CNS Event Agent	4-8
Enabling the CNS Configuration Agent	4-9
Enabling an Initial Configuration	4-9
Enabling a Partial Configuration	4-11
Displaying CNS Configuration	4-12

CHAPTER 5

Administering the Switch	5-1
Managing the System Time and Date	5-1
Understanding the System Clock	5-2
Understanding Network Time Protocol	5-2
Configuring NTP	5-4
Default NTP Configuration	5-4
Configuring NTP Authentication	5-5
Configuring NTP Associations	5-6
Configuring NTP Broadcast Service	5-7
Configuring NTP Access Restrictions	5-8
Configuring the Source IP Address for NTP Packets	5-10
Displaying the NTP Configuration	5-11
Configuring Time and Date Manually	5-11
Setting the System Clock	5-11
Displaying the Time and Date Configuration	5-12
Configuring the Time Zone	5-12
Configuring Summer Time (Daylight Saving Time)	5-13

Configuring a System Name and Prompt	5-15
Default System Name and Prompt Configuration	5-15
Configuring a System Name	5-15
Configuring a System Prompt	5-16
Understanding DNS	5-16
Default DNS Configuration	5-17
Setting Up DNS	5-17
Displaying the DNS Configuration	5-18
Creating a Banner	5-18
Default Banner Configuration	5-18
Configuring a Message-of-the-Day Login Banner	5-19
Configuring a Login Banner	5-20
Managing the MAC Address Table	5-20
Building the Address Table	5-21
MAC Addresses and VLANs	5-21
Default MAC Address Table Configuration	5-21
Changing the Address Aging Time	5-22
Removing Dynamic Address Entries	5-22
Configuring MAC Address Notification Traps	5-23
Adding and Removing Static Address Entries	5-24
Displaying Address Table Entries	5-25
Managing the ARP Table	5-26

CHAPTER 6

Configuring SDM Templates	6-1
Understanding the SDM Templates	6-1
Configuring the Switch SDM Template	6-2
Default SDM Template	6-2
SDM Template Configuration Guidelines	6-3
Setting the SDM Template	6-3
Displaying the SDM Templates	6-4

CHAPTER 7

Configuring Switch-Based Authentication	7-1
Preventing Unauthorized Access to Your Switch	7-1
Protecting Access to Privileged EXEC Commands	7-2
Default Password and Privilege Level Configuration	7-2
Setting or Changing a Static Enable Password	7-3
Protecting Enable and Enable Secret Passwords with Encryption	7-3
Disabling Password Recovery	7-5
Setting a Telnet Password for a Terminal Line	7-6

Configuring Username and Password Pairs	7-6
Configuring Multiple Privilege Levels	7-8
Setting the Privilege Level for a Command	7-8
Changing the Default Privilege Level for Lines	7-9
Logging into and Exiting a Privilege Level	7-10
Controlling Switch Access with TACACS+	7-10
Understanding TACACS+	7-10
TACACS+ Operation	7-12
Configuring TACACS+	7-12
Default TACACS+ Configuration	7-13
Identifying the TACACS+ Server Host and Setting the Authentication Key	7-13
Configuring TACACS+ Login Authentication	7-14
Configuring TACACS+ Authorization for Privileged EXEC Access and Network Services	7-16
Starting TACACS+ Accounting	7-16
Displaying the TACACS+ Configuration	7-17
Controlling Switch Access with RADIUS	7-17
Understanding RADIUS	7-18
RADIUS Operation	7-19
Configuring RADIUS	7-20
Default RADIUS Configuration	7-20
Identifying the RADIUS Server Host	7-20
Configuring RADIUS Login Authentication	7-23
Defining AAA Server Groups	7-25
Configuring RADIUS Authorization for User Privileged Access and Network Services	7-27
Starting RADIUS Accounting	7-28
Configuring Settings for All RADIUS Servers	7-28
Configuring the Switch to Use Vendor-Specific RADIUS Attributes	7-29
Configuring the Switch for Vendor-Proprietary RADIUS Server Communication	7-30
Displaying the RADIUS Configuration	7-31
Controlling Switch Access with Kerberos	7-31
Understanding Kerberos	7-32
Kerberos Operation	7-34
Authenticating to a Boundary Switch	7-34
Obtaining a TGT from a KDC	7-35
Authenticating to Network Services	7-35
Configuring Kerberos	7-35
Configuring the Switch for Local Authentication and Authorization	7-36

Configuring the Switch for Secure Shell 7-37

Understanding SSH 7-37

Configuring SSH 7-37

CHAPTER 8

Configuring 802.1x Port-Based Authentication 8-1

Understanding 802.1x Port-Based Authentication 8-1

Device Roles 8-2

Authentication Initiation and Message Exchange 8-3

Ports in Authorized and Unauthorized States 8-4

Supported Topologies 8-5

Using 802.1x with Port Security 8-6

Using 802.1x with Voice VLAN Ports 8-6

Using 802.1x with VLAN Assignment 8-7

Using 802.1x with Guest VLAN 8-8

Using 802.1x with Per-User ACLs 8-8

Configuring 802.1x Authentication 8-9

Default 802.1x Configuration 8-10

802.1x Configuration Guidelines 8-11

Configuring 802.1x Authentication 8-11

Configuring the Switch-to-RADIUS-Server Communication 8-13

Configuring Periodic Re-Authentication 8-14

Manually Re-Authenticating a Client Connected to a Port 8-14

Changing the Quiet Period 8-15

Changing the Switch-to-Client Retransmission Time 8-15

Setting the Switch-to-Client Frame-Retransmission Number 8-16

Configuring the Host Mode 8-17

Configuring a Guest VLAN 8-18

Resetting the 802.1x Configuration to the Default Values 8-18

Displaying 802.1x Statistics and Status 8-19

CHAPTER 9

Configuring Interface Characteristics 9-1

Understanding Interface Types 9-1

Port-Based VLANs 9-2

Switch Ports 9-2

Access Ports 9-2

Trunk Ports 9-3

Tunnel Ports 9-3

Routed Ports 9-4

Switch Virtual Interfaces 9-4

EtherChannel Port Groups	9-5
Connecting Interfaces	9-5
Using Interface Configuration Mode	9-6
Procedures for Configuring Interfaces	9-7
Configuring a Range of Interfaces	9-8
Configuring and Using Interface Range Macros	9-9
Configuring Ethernet Interfaces	9-11
Default Ethernet Interface Configuration	9-11
Configuring Interface Speed and Duplex Mode	9-12
Configuration Guidelines	9-13
Setting the Interface Speed and Duplex Parameters	9-13
Configuring IEEE 802.3z Flow Control	9-15
Configuring Auto-MDIX on a Port	9-16
Adding a Description for an Interface	9-17
Configuring Layer 3 Interfaces	9-18
Configuring the System MTU	9-19
Monitoring and Maintaining the Interfaces	9-20
Monitoring Interface Status	9-20
Clearing and Resetting Interfaces and Counters	9-21
Shutting Down and Restarting the Interface	9-22

CHAPTER 10

Configuring VLANs 10-1

Understanding VLANs	10-1
Supported VLANs	10-3
VLAN Port Membership Modes	10-3
Configuring Normal-Range VLANs	10-4
Token Ring VLANs	10-6
Normal-Range VLAN Configuration Guidelines	10-6
VLAN Configuration Mode Options	10-7
VLAN Configuration in config-vlan Mode	10-7
VLAN Configuration in VLAN Database Configuration Mode	10-7
Saving VLAN Configuration	10-7
Default Ethernet VLAN Configuration	10-8
Creating or Modifying an Ethernet VLAN	10-9
Deleting a VLAN	10-11
Assigning Static-Access Ports to a VLAN	10-11

Configuring Extended-Range VLANs	10-12
Default VLAN Configuration	10-13
Extended-Range VLAN Configuration Guidelines	10-13
Creating an Extended-Range VLAN	10-14
Creating an Extended-Range VLAN with an Internal VLAN ID	10-15
Displaying VLANs	10-16
Configuring VLAN Trunks	10-16
Trunking Overview	10-16
Encapsulation Types	10-18
802.1Q Configuration Considerations	10-19
Default Layer 2 Ethernet Interface VLAN Configuration	10-19
Configuring an Ethernet Interface as a Trunk Port	10-20
Interaction with Other Features	10-20
Configuring a Trunk Port	10-21
Defining the Allowed VLANs on a Trunk	10-22
Changing the Pruning-Eligible List	10-23
Configuring the Native VLAN for Untagged Traffic	10-24
Configuring Trunk Ports for Load Sharing	10-24
Load Sharing Using STP Port Priorities	10-25
Load Sharing Using STP Path Cost	10-26
Configuring VMPS	10-28
Understanding VMPS	10-28
Dynamic-Access Port VLAN Membership	10-29
Default VMPS Client Configuration	10-29
VMPS Configuration Guidelines	10-29
Configuring the VMPS Client	10-30
Entering the IP Address of the VMPS	10-30
Configuring Dynamic-Access Ports on VMPS Clients	10-31
Reconfirming VLAN Memberships	10-31
Changing the Reconfirmation Interval	10-32
Changing the Retry Count	10-32
Monitoring the VMPS	10-33
Troubleshooting Dynamic-Access Port VLAN Membership	10-33
VMPS Configuration Example	10-34

CHAPTER 11

Configuring VTP 11-1

- Understanding VTP 11-1
 - The VTP Domain 11-2
 - VTP Modes 11-3
 - VTP Advertisements 11-3
 - VTP Version 2 11-4
 - VTP Pruning 11-4
- Configuring VTP 11-6
 - Default VTP Configuration 11-6
 - VTP Configuration Options 11-7
 - VTP Configuration in Global Configuration Mode 11-7
 - VTP Configuration in VLAN Database Configuration Mode 11-7
 - VTP Configuration Guidelines 11-8
 - Domain Names 11-8
 - Passwords 11-8
 - VTP Version 11-8
 - Configuration Requirements 11-9
 - Configuring a VTP Server 11-9
 - Configuring a VTP Client 11-11
 - Disabling VTP (VTP Transparent Mode) 11-12
 - Enabling VTP Version 2 11-13
 - Enabling VTP Pruning 11-14
 - Adding a VTP Client Switch to a VTP Domain 11-15
- Monitoring VTP 11-16

CHAPTER 12

Configuring Voice VLAN 12-1

- Understanding Voice VLAN 12-1
 - Cisco IP Phone Voice Traffic 12-2
 - Cisco IP Phone Data Traffic 12-2
- Configuring Voice VLAN 12-3
 - Default Voice VLAN Configuration 12-3
 - Voice VLAN Configuration Guidelines 12-3
 - Configuring a Port Connected to a Cisco 7960 IP Phone 12-4
 - Configuring IP Phone Voice Traffic 12-4
 - Configuring the Priority of Incoming Data Frames 12-5
- Displaying Voice VLAN 12-6

CHAPTER 13

Configuring IEEE 802.1Q and Layer 2 Protocol Tunneling 13-1

- Understanding 802.1Q Tunneling 13-2
- Configuring 802.1Q Tunneling 13-4
 - Default 802.1Q Tunneling Configuration 13-4
 - 802.1Q Tunneling Configuration Guidelines 13-4
 - Native VLANs 13-4
 - System MTU 13-5
 - 802.1Q Tunneling and Other Features 13-6
 - Configuring an 802.1Q Tunneling Port 13-6
- Configuring VLAN Mapping 13-7
 - Default VLAN Mapping Configuration 13-8
 - Mapping Customer VLANs to Service-Provider VLANs 13-8
 - Mapping Customer 802.1Q Traffic with VLAN IDs 13-9
- Understanding Layer 2 Protocol Tunneling 13-10
- Configuring Layer 2 Protocol Tunneling 13-12
 - Default Layer 2 Protocol Tunneling Configuration 13-13
 - Layer 2 Protocol Tunneling Configuration Guidelines 13-13
 - Configuring Layer 2 Tunneling 13-14
- Monitoring and Maintaining Tunneling and Mapping Status 13-16

CHAPTER 14

Configuring STP 14-1

- Understanding Spanning-Tree Features 14-1
- STP Overview 14-2
- Spanning-Tree Topology and BPDUs 14-3
- Bridge ID, Switch Priority, and Extended System ID 14-4
- Spanning-Tree Interface States 14-4
 - Blocking State 14-5
 - Listening State 14-6
 - Learning State 14-6
 - Forwarding State 14-6
 - Disabled State 14-7
- How a Switch or Port Becomes the Root Switch or Root Port 14-7
- Spanning Tree and Redundant Connectivity 14-8
- Spanning-Tree Address Management 14-8
- Accelerated Aging to Retain Connectivity 14-8
- Spanning-Tree Modes and Protocols 14-9
- Supported Spanning-Tree Instances 14-10
- Spanning-Tree Interoperability and Backward Compatibility 14-10
- STP and IEEE 802.1Q Trunks 14-10

VLAN-Bridge Spanning Tree	14-11
Configuring Spanning-Tree Features	14-11
Default Spanning-Tree Configuration	14-11
Spanning-Tree Configuration Guidelines	14-12
Changing the Spanning-Tree Mode	14-13
Disabling Spanning Tree	14-14
Configuring the Root Switch	14-14
Configuring a Secondary Root Switch	14-16
Configuring Port Priority	14-17
Configuring Path Cost	14-18
Configuring the Switch Priority of a VLAN	14-19
Configuring Spanning-Tree Timers	14-20
Configuring the Hello Time	14-20
Configuring the Forwarding-Delay Time for a VLAN	14-21
Configuring the Maximum-Aging Time for a VLAN	14-21
Displaying the Spanning-Tree Status	14-22

CHAPTER 15

Configuring MSTP 15-1

Understanding MSTP	15-2
Multiple Spanning-Tree Regions	15-2
IST, CIST, and CST	15-3
Operations Within an MST Region	15-3
Operations Between MST Regions	15-4
Hop Count	15-5
Boundary Ports	15-5
Interoperability with 802.1D STP	15-5
Understanding RSTP	15-6
Port Roles and the Active Topology	15-6
Rapid Convergence	15-7
Synchronization of Port Roles	15-8
Bridge Protocol Data Unit Format and Processing	15-9
Processing Superior BPDU Information	15-10
Processing Inferior BPDU Information	15-10
Topology Changes	15-10
Configuring MSTP Features	15-11
Default MSTP Configuration	15-12
MSTP Configuration Guidelines	15-12
Specifying the MST Region Configuration and Enabling MSTP	15-13
Configuring the Root Switch	15-14

Configuring a Secondary Root Switch	15-16
Configuring Port Priority	15-17
Configuring Path Cost	15-18
Configuring the Switch Priority	15-19
Configuring the Hello Time	15-19
Configuring the Forwarding-Delay Time	15-20
Configuring the Maximum-Aging Time	15-21
Configuring the Maximum-Hop Count	15-21
Specifying the Link Type to Ensure Rapid Transitions	15-22
Restarting the Protocol Migration Process	15-22
Displaying the MST Configuration and Status	15-23

CHAPTER 16

Configuring Optional Spanning-Tree Features 16-1

Understanding Optional Spanning-Tree Features	16-1
Understanding Port Fast	16-2
Understanding BPDU Guard	16-3
Understanding BPDU Filtering	16-3
Understanding UplinkFast	16-4
Understanding BackboneFast	16-5
Understanding Root Guard	16-7
Understanding Loop Guard	16-8
Configuring Optional Spanning-Tree Features	16-9
Default Optional Spanning-Tree Configuration	16-9
Optional Spanning-Tree Configuration Guidelines	16-9
Enabling Port Fast	16-10
Enabling BPDU Guard	16-11
Enabling BPDU Filtering	16-12
Enabling UplinkFast for Use with Redundant Links	16-13
Enabling BackboneFast	16-14
Enabling Root Guard	16-14
Enabling Loop Guard	16-15
Displaying the Spanning-Tree Status	16-16

CHAPTER 17

Configuring IGMP Snooping and MVR 17-1

Understanding IGMP Snooping	17-2
Joining a Multicast Group	17-2
Leaving a Multicast Group	17-4
Immediate-Leave Processing	17-5

Configuring IGMP Snooping	17-5
Default IGMP Snooping Configuration	17-5
Enabling or Disabling IGMP Snooping	17-6
Setting the Snooping Method	17-6
Configuring a Multicast Router Port	17-8
Configuring a Host Statically to Join a Group	17-9
Enabling IGMP Immediate-Leave Processing	17-10
Displaying IGMP Snooping Information	17-10
Understanding Multicast VLAN Registration	17-12
Using MVR in a Multicast Television Application	17-13
Configuring MVR	17-14
Default MVR Configuration	17-14
MVR Configuration Guidelines and Limitations	17-15
Configuring MVR Global Parameters	17-15
Configuring MVR Interfaces	17-17
Displaying MVR Information	17-18
Configuring IGMP Filtering	17-19
Default IGMP Filtering Configuration	17-19
Configuring IGMP Profiles	17-19
Applying IGMP Profiles	17-20
Setting the Maximum Number of IGMP Groups	17-21
Displaying IGMP Filtering Configuration	17-22

CHAPTER 18

Configuring Port-Based Traffic Control	18-1
Configuring Storm Control	18-1
Understanding Storm Control	18-1
Default Storm Control Configuration	18-3
Enabling Storm Control	18-3
Configuring Protected Ports	18-4
Default Protected Port Configuration	18-4
Protected Port Configuration Guidelines	18-5
Configuring a Protected Port	18-5
Configuring Port Blocking	18-5
Default Port Blocking Configuration	18-5
Blocking Flooded Traffic on an Interface	18-6

Configuring Port Security	18-6
Understanding Port Security	18-7
Secure MAC Addresses	18-7
Security Violations	18-8
Default Port Security Configuration	18-9
Configuration Guidelines	18-9
Enabling and Configuring Port Security	18-10
Enabling and Configuring Port Security Aging	18-13
Displaying Port-Based Traffic Control Settings	18-14

CHAPTER 19

Configuring CDP	19-1
Understanding CDP	19-1
Configuring CDP	19-2
Default CDP Configuration	19-2
Configuring the CDP Characteristics	19-2
Disabling and Enabling CDP	19-3
Disabling and Enabling CDP on an Interface	19-4
Monitoring and Maintaining CDP	19-5

CHAPTER 20

Configuring UDLD	20-1
Understanding UDLD	20-1
Modes of Operation	20-1
Methods to Detect Unidirectional Links	20-2
Configuring UDLD	20-4
Default UDLD Configuration	20-4
UDLD Configuration Guidelines	20-4
Enabling UDLD Globally	20-5
Enabling UDLD on an Interface	20-5
Resetting an Interface Disabled by UDLD	20-6
Displaying UDLD Status	20-6

CHAPTER 21

Configuring SPAN and RSPAN	21-1
Understanding SPAN and RSPAN	21-1
Local SPAN	21-2
Remote SPAN	21-2

SPAN and RSPAN Concepts and Terminology	21-3
SPAN Sessions	21-3
Monitored Traffic	21-4
Source Ports	21-5
Source VLANs	21-6
VLAN Filtering	21-6
Destination Port	21-7
RSPAN VLAN	21-8
SPAN and RSPAN Interaction with Other Features	21-8
Configuring SPAN and RSPAN	21-9
Default SPAN and RSPAN Configuration	21-9
Configuring Local SPAN	21-10
SPAN Configuration Guidelines	21-10
Creating a Local SPAN Session	21-11
Creating a Local SPAN Session and Configuring Ingress Traffic	21-13
Specifying VLANs to Filter	21-15
Configuring RSPAN	21-16
RSPAN Configuration Guidelines	21-16
Configuring a VLAN as an RSPAN VLAN	21-17
Creating an RSPAN Source Session	21-18
Creating an RSPAN Destination Session	21-19
Creating an RSPAN Destination Session and Configuring Ingress Traffic	21-20
Specifying VLANs to Filter	21-22
Displaying SPAN and RSPAN Status	21-23

CHAPTER 22

Configuring RMON	22-1
Understanding RMON	22-1
Configuring RMON	22-2
Default RMON Configuration	22-3
Configuring RMON Alarms and Events	22-3
Collecting Group History Statistics on an Interface	22-5
Collecting Group Ethernet Statistics on an Interface	22-6
Displaying RMON Status	22-6

CHAPTER 23

Configuring System Message Logging	23-1
Understanding System Message Logging	23-1
Configuring System Message Logging	23-2
System Log Message Format	23-2
Default System Message Logging Configuration	23-3

Disabling Message Logging	23-4
Setting the Message Display Destination Device	23-4
Synchronizing Log Messages	23-5
Enabling and Disabling Timestamps on Log Messages	23-7
Enabling and Disabling Sequence Numbers in Log Messages	23-7
Defining the Message Severity Level	23-8
Limiting Syslog Messages Sent to the History Table and to SNMP	23-9
Configuring UNIX Syslog Servers	23-10
Logging Messages to a UNIX Syslog Daemon	23-10
Configuring the UNIX System Logging Facility	23-11
Displaying the Logging Configuration	23-12

CHAPTER 24

Configuring SNMP 24-1

Understanding SNMP	24-1
SNMP Versions	24-2
SNMP Manager Functions	24-3
SNMP Agent Functions	24-4
SNMP Community Strings	24-4
Using SNMP to Access MIB Variables	24-4
SNMP Notifications	24-5
Configuring SNMP	24-5
Default SNMP Configuration	24-6
SNMP Configuration Guidelines	24-6
Disabling the SNMP Agent	24-7
Configuring Community Strings	24-7
Configuring SNMP Groups and Users	24-8
Configuring SNMP Notifications	24-10
Setting the Agent Contact and Location Information	24-13
Limiting TFTP Servers Used Through SNMP	24-13
SNMP Examples	24-14
Displaying SNMP Status	24-15

CHAPTER 25

Configuring Network Security with ACLs 25-1

Understanding ACLs	25-1
Supported ACLs	25-2
Router ACLs	25-3
Port ACLs	25-3
VLAN Maps	25-4
Handling Fragmented and Unfragmented Traffic	25-5

Configuring IP ACLs	25-6
Creating Standard and Extended IP ACLs	25-6
Access List Numbers	25-7
Creating a Numbered Standard ACL	25-8
Creating a Numbered Extended ACL	25-9
Creating Named Standard and Extended ACLs	25-13
Using Time Ranges with ACLs	25-15
Including Comments in ACLs	25-17
Applying an IP ACL to a Terminal Line	25-17
Applying an IP ACL to an Interface	25-18
Hardware and Software Treatment of IP ACLs	25-19
IP ACL Configuration Examples	25-20
Numbered ACLs	25-21
Extended ACLs	25-21
Named ACLs	25-22
Time Range Applied to an IP ACL	25-23
Commented IP ACL Entries	25-23
ACL Logging	25-24
Creating Named MAC Extended ACLs	25-25
Applying a MAC ACL to a Layer 2 Interface	25-26
Configuring VLAN Maps	25-27
VLAN Map Configuration Guidelines	25-28
Creating a VLAN Map	25-28
Examples of ACLs and VLAN Maps	25-29
Applying a VLAN Map to a VLAN	25-31
Using VLAN Maps in Your Network	25-31
Wiring Closet Configuration	25-32
Denying Access to a Server on Another VLAN	25-33
Using VLAN Maps with Router ACLs	25-34
Guidelines	25-34
Examples of Router ACLs and VLAN Maps Applied to VLANs	25-35
ACLs and Switched Packets	25-35
ACLs and Bridged Packets	25-36
ACLs and Routed Packets	25-36
ACLs and Multicast Packets	25-37
Displaying ACL Configuration	25-38

Configuring QoS	26-1
Understanding QoS	26-2
Basic QoS Model	26-4
Ingress Classification	26-6
Ingress Classification Based on QoS ACLs	26-8
Ingress Classification Based on Traffic Classes and Traffic Policies	26-8
Ingress Policing and Marking	26-9
Mapping Tables	26-11
Queueing and Scheduling Overview	26-12
Weighted Tail Drop	26-13
SRR Shaping and Sharing	26-14
Queueing and Scheduling of Ingress Queues	26-15
Queueing and Scheduling of Egress Queue-Sets	26-17
Understanding Hierarchical QoS	26-19
Hierarchical Levels	26-20
Egress Classification Based on Traffic Classes and Traffic Policies	26-23
Egress Policing and Marking	26-24
Queueing and Scheduling of Hierarchical Queues	26-26
Hierarchical Queues	26-27
Congestion-Management and Congestion-Avoidance Features	26-27
Configuring Auto-QoS	26-29
Generated Auto-QoS Configuration	26-30
Effects of Auto-QoS on the Configuration	26-33
Auto-QoS Configuration Guidelines	26-33
Enabling Auto-QoS for VoIP	26-34
Auto-QoS Configuration Example	26-35
Displaying Auto-QoS Information	26-37
Configuring Standard QoS	26-37
Default Standard QoS Configuration	26-38
Default Ingress Queue Configuration	26-38
Default Egress Queue-Set Configuration	26-39
Default Mapping Table Configuration	26-39
Standard QoS Configuration Guidelines	26-40
Packet Modification	26-41
Enabling QoS Globally	26-42
Configuring Ingress Classification by Using Port Trust States	26-42
Configuring the Trust State on Ports Within the QoS Domain	26-42
Configuring the CoS Value for an Interface	26-45
Configuring a Trusted Boundary to Ensure Port Security	26-46

Configuring the DSCP Trust State on a Port Bordering Another QoS Domain	26-47
Configuring an Ingress QoS Policy	26-48
Classifying Ingress Traffic by Using ACLs	26-49
Classifying Ingress Traffic by Using Class Maps	26-52
Classifying, Policing, and Marking Ingress Traffic by Using Policy Maps	26-54
Classifying, Policing, and Marking Ingress Traffic by Using Aggregate Policers	26-57
Configuring DSCP Maps	26-59
Configuring the CoS-to-DSCP Map	26-59
Configuring the IP-Precedence-to-DSCP Map	26-60
Configuring the Policed-DSCP Map	26-61
Configuring the DSCP-to-CoS Map	26-62
Configuring the DSCP-to-DSCP-Mutation Map	26-63
Configuring Ingress Queue Characteristics	26-64
Mapping DSCP or CoS Values to an Ingress Queue and Setting WTD Thresholds	26-65
Allocating Buffer Space Between the Ingress Queues	26-66
Allocating Bandwidth Between the Ingress Queues	26-67
Configuring the Ingress Priority Queue	26-68
Configuring Egress Queue-Set Characteristics	26-69
Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set	26-69
Mapping DSCP or CoS Values to an Egress Queue-Set and to a Threshold ID	26-71
Configuring SRR Shaped Weights on an Egress Queue-Set	26-72
Configuring SRR Shared Weights on an Egress Queue-Set	26-73
Limiting the Egress Bandwidth on a Queue-Set	26-74
Displaying Standard QoS Information	26-75
Configuring Hierarchical QoS	26-76
Default Hierarchical QoS Configuration	26-76
Hierarchical QoS Configuration Guidelines	26-76
Configuring an Egress Hierarchical QoS Policy	26-77
Classifying Egress Traffic by Using Class Maps	26-78
Configuring an Egress Two-Rate Traffic Policer	26-80
Configuring Class-Based Packet Marking in an Egress Traffic Policy	26-84
Configuring CBWFQ and Tail Drop	26-86
Configuring CBWFQ and DSCP-Based WRED	26-89
Configuring CBWFQ and IP Precedence-Based WRED	26-93
Enabling LLQ	26-97
Configuring Shaping	26-99
Displaying Hierarchical QoS Information	26-101

CHAPTER 27

Configuring EtherChannels 27-1

- Understanding EtherChannels 27-1
 - EtherChannel Overview 27-2
 - Port-Channel Interfaces 27-3
 - Port Aggregation Protocol 27-4
 - PAgP Modes 27-4
 - PAgP Interaction with Other Features 27-5
 - Link Aggregation Control Protocol 27-5
 - LACP Modes 27-6
 - LACP Interaction with Other Features 27-6
 - Load Balancing and Forwarding Methods 27-6
- Configuring EtherChannels 27-8
 - Default EtherChannel Configuration 27-9
 - EtherChannel Configuration Guidelines 27-9
 - Configuring Layer 2 EtherChannels 27-10
 - Configuring Layer 3 EtherChannels 27-12
 - Creating Port-Channel Logical Interfaces 27-12
 - Configuring the Physical Interfaces 27-13
 - Configuring EtherChannel Load Balancing 27-15
 - Configuring the PAgP Learn Method and Priority 27-16
 - Configuring LACP Hot-Standby Ports 27-17
 - Configuring the LACP System Priority 27-18
 - Configuring the LACP Port Priority 27-18
- Displaying EtherChannel, PAgP, and LACP Status 27-19

CHAPTER 28

Configuring IP Unicast Routing 28-1

- Understanding IP Routing 28-2
- Steps for Configuring Routing 28-3
- Configuring IP Addressing 28-4
 - Default Addressing Configuration 28-4
 - Assigning IP Addresses to Network Interfaces 28-5
 - Use of Subnet Zero 28-6
 - Classless Routing 28-6
- Configuring Address Resolution Methods 28-7
 - Define a Static ARP Cache 28-8
 - Set ARP Encapsulation 28-9
 - Enable Proxy ARP 28-10

Routing Assistance When IP Routing is Disabled	28-10
Proxy ARP	28-10
Default Gateway	28-10
ICMP Router Discovery Protocol (IRDP)	28-11
Configuring Broadcast Packet Handling	28-12
Enabling Directed Broadcast-to-Physical Broadcast Translation	28-13
Forwarding UDP Broadcast Packets and Protocols	28-14
Establishing an IP Broadcast Address	28-15
Flooding IP Broadcasts	28-15
Monitoring and Maintaining IP Addressing	28-16
Enabling IP Unicast Routing	28-17
Configuring RIP	28-18
Default RIP Configuration	28-18
Configuring Basic RIP Parameters	28-19
Configuring RIP Authentication	28-21
Configuring Summary Addresses and Split Horizon	28-21
Configuring IGRP	28-23
Default IGRP Configuration	28-24
Understanding Load Balancing and Traffic Distribution Control	28-24
Configuring Basic IGRP Parameters	28-25
Configuring Split Horizon	28-27
Configuring OSPF	28-28
Default OSPF Configuration	28-29
Configuring Basic OSPF Parameters	28-30
Configuring OSPF Interfaces	28-31
Configuring OSPF Area Parameters	28-32
Configuring Other OSPF Parameters	28-33
Changing LSA Group Pacing	28-35
Configuring a Loopback Interface	28-35
Monitoring OSPF	28-36
Configuring EIGRP	28-37
Default EIGRP Configuration	28-38
Configuring Basic EIGRP Parameters	28-39
Configuring EIGRP Interfaces	28-40
Configuring EIGRP Route Authentication	28-41
Monitoring and Maintaining EIGRP	28-42

Configuring BGP	28-43
Default BGP Configuration	28-45
Enabling BGP Routing	28-47
Managing Routing Policy Changes	28-49
Configuring BGP Decision Attributes	28-50
Configuring BGP Filtering with Route Maps	28-52
Configuring BGP Filtering by Neighbor	28-53
Configuring Prefix Lists for BGP Filtering	28-54
Configuring BGP Community Filtering	28-55
Configuring BGP Neighbors and Peer Groups	28-56
Configuring Aggregate Addresses	28-58
Configuring Routing Domain Confederations	28-59
Configuring BGP Route Reflectors	28-60
Configuring Route Dampening	28-61
Monitoring and Maintaining BGP	28-62
Configuring ISO CLNS Routing	28-63
Configuring ISO IGRP Dynamic Routing	28-63
Default ISO IGRP Configuration	28-64
Enabling ISO IGRP	28-64
Configuring Optional ISO IGRP Parameters	28-65
Configuring IS-IS Dynamic Routing	28-66
Default IS-IS Configuration	28-67
Enabling IS-IS Routing	28-68
Configuring IS-IS Global Parameters	28-70
Configuring IS-IS Interface Parameters	28-72
Monitoring and Maintaining ISO IGRP and IS-IS	28-74
Configuring Multi-VRF CE	28-76
Understanding Multi-VRF CE	28-76
Default Multi-VRF CE Configuration	28-78
Multi-VRF CE Configuration Guidelines	28-79
Configuring VRFs	28-80
Configuring a VPN Routing Session	28-80
Configuring BGP PE to CE Routing Sessions	28-81
Multi-VRF CE Configuration Example	28-82
Displaying Multi-VRF CE Status	28-86
Configuring Protocol-Independent Features	28-86
Configuring Cisco Express Forwarding	28-86
Configuring the Number of Equal-Cost Routing Paths	28-87
Configuring Static Unicast Routes	28-88

Specifying Default Routes and Networks	28-89
Using Route Maps to Redistribute Routing Information	28-90
Configuring Policy-Based Routing	28-94
PBR Configuration Guidelines	28-94
Enabling PBR	28-95
Filtering Routing Information	28-96
Setting Passive Interfaces	28-97
Controlling Advertising and Processing in Routing Updates	28-97
Filtering Sources of Routing Information	28-98
Managing Authentication Keys	28-99
Monitoring and Maintaining the IP Network	28-100

CHAPTER 29

Configuring HSRP	29-1
Understanding HSRP	29-1
Configuring HSRP	29-3
Default HSRP Configuration	29-4
HSRP Configuration Guidelines	29-4
Enabling HSRP	29-5
Configuring HSRP Group Attributes	29-6
Configuring HSRP Priority	29-6
Configuring HSRP Authentication and Timers	29-8
Displaying HSRP Configurations	29-10

CHAPTER 30

Configuring MPLS and EoMPLS	30-1
Understanding MPLS Services	30-1
Understanding MPLS VPNs	30-2
VPN Benefits	30-3
Distribution of VPN Routing Information	30-5
Configuring MPLS VPNs	30-6
Default MPLS Configuration	30-6
MPLS VPN Configuration Guidelines	30-6
Enabling MPLS	30-7
Defining VPNs	30-8
Configuring BGP Routing Sessions	30-9
Configuring PE-to-PE Routing Sessions	30-9
Configuring BGP PE-to-CE Routing Sessions	30-10
Configuring RIP PE-to-CE Routing Sessions	30-10
Configuring Static Route PE-to-CE Routing Sessions	30-11
Packet Flow in an MPLS VPN	30-11

Understanding EoMPLS	30-12
Interaction with Other Features	30-13
EoMPLS and 802.1Q Tunneling	30-13
EoMPLS and Layer 2 Tunneling	30-14
EoMPLS and QoS	30-14
EoMPLS Limitations	30-14
Enabling EoMPLS	30-15
Default EoMPLS Configuration	30-15
EoMPLS Configuration Guidelines	30-15
Configuring EoMPLS	30-16
Packet Flow in an EoMPLS Network	30-17
Configuring MPLS and EoMPLS QoS	30-18
Understanding MPLS QoS	30-18
Enabling MPLS and EoMPLS QoS	30-20
Default MPLS and EoMPLS QoS Configuration	30-20
Setting the Priority of Packets with Experimental Bits	30-20
Monitoring and Maintaining MPLS and EoMPLS	30-22

 CHAPTER 31

Configuring IP Multicast Routing	31-1
Understanding Cisco's Implementation of IP Multicast Routing	31-2
Understanding IGMP	31-2
IGMP Version 1	31-3
IGMP Version 2	31-3
Understanding PIM	31-3
PIM Versions	31-4
PIM Modes	31-4
Auto-RP	31-5
Bootstrap Router	31-5
Multicast Forwarding and Reverse Path Check	31-6
Understanding DVMRP	31-7
Understanding CGMP	31-7
Configuring IP Multicast Routing	31-8
Default Multicast Routing Configuration	31-8
Multicast Routing Configuration Guidelines	31-8
PIMv1 and PIMv2 Interoperability	31-9
Auto-RP and BSR Configuration Guidelines	31-9
Configuring Basic Multicast Routing	31-10

Configuring a Rendezvous Point	31-11
Manually Assigning an RP to Multicast Groups	31-12
Configuring Auto-RP	31-13
Configuring PIMv2 BSR	31-17
Using Auto-RP and a BSR	31-21
Monitoring the RP Mapping Information	31-22
Troubleshooting PIMv1 and PIMv2 Interoperability Problems	31-22
Configuring Advanced PIM Features	31-22
Understanding PIM Shared Tree and Source Tree	31-22
Delaying the Use of PIM Shortest-Path Tree	31-24
Modifying the PIM Router-Query Message Interval	31-25
Configuring Optional IGMP Features	31-26
Default IGMP Configuration	31-26
Configuring the Switch as a Member of a Group	31-26
Controlling Access to IP Multicast Groups	31-27
Changing the IGMP Version	31-28
Modifying the IGMP Host-Query Message Interval	31-29
Changing the IGMP Query Timeout for IGMPv2	31-30
Changing the Maximum Query Response Time for IGMPv2	31-30
Configuring the Switch as a Statically Connected Member	31-31
Configuring Optional Multicast Routing Features	31-31
Enabling CGMP Server Support	31-32
Configuring sdr Listener Support	31-33
Enabling sdr Listener Support	31-33
Limiting How Long an sdr Cache Entry Exists	31-34
Configuring an IP Multicast Boundary	31-34
Configuring Basic DVMRP Interoperability Features	31-36
Configuring DVMRP Interoperability	31-36
Configuring a DVMRP Tunnel	31-38
Advertising Network 0.0.0.0 to DVMRP Neighbors	31-40
Responding to mrinfo Requests	31-41
Configuring Advanced DVMRP Interoperability Features	31-41
Enabling DVMRP Unicast Routing	31-42
Rejecting a DVMRP Nonpruning Neighbor	31-43
Controlling Route Exchanges	31-45
Limiting the Number of DVMRP Routes Advertised	31-45
Changing the DVMRP Route Threshold	31-45
Configuring a DVMRP Summary Address	31-46
Disabling DVMRP Autosummarization	31-48

Adding a Metric Offset to the DVMRP Route	31-48
Monitoring and Maintaining IP Multicast Routing	31-49
Clearing Caches, Tables, and Databases	31-49
Displaying System and Network Statistics	31-50
Monitoring IP Multicast Routing	31-51

CHAPTER 32

Configuring MSDP	32-1
Understanding MSDP	32-1
MSDP Operation	32-2
MSDP Benefits	32-3
Configuring MSDP	32-4
Default MSDP Configuration	32-4
Configuring a Default MSDP Peer	32-4
Caching Source-Active State	32-6
Requesting Source Information from an MSDP Peer	32-8
Controlling Source Information that Your Switch Originates	32-8
Redistributing Sources	32-9
Filtering Source-Active Request Messages	32-11
Controlling Source Information that Your Switch Forwards	32-12
Using a Filter	32-12
Using TTL to Limit the Multicast Data Sent in SA Messages	32-14
Controlling Source Information that Your Switch Receives	32-14
Configuring an MSDP Mesh Group	32-16
Shutting Down an MSDP Peer	32-16
Including a Bordering PIM Dense-Mode Region in MSDP	32-17
Configuring an Originating Address other than the RP Address	32-18
Monitoring and Maintaining MSDP	32-19

CHAPTER 33

Configuring Fallback Bridging	33-1
Understanding Fallback Bridging	33-1
Configuring Fallback Bridging	33-2
Default Fallback Bridging Configuration	33-3
Fallback Bridging Configuration Guidelines	33-3
Creating a Bridge Group	33-3

Adjusting Spanning-Tree Parameters	33-5
Changing the VLAN-Bridge Spanning-Tree Priority	33-5
Changing the Interface Priority	33-6
Assigning a Path Cost	33-7
Adjusting BPDU Intervals	33-8
Disabling the Spanning Tree on an Interface	33-10
Monitoring and Maintaining Fallback Bridging	33-10

CHAPTER 34

Troubleshooting 34-1

Recovering from Corrupted Software By Using the XMODEM Protocol	34-2
Recovering from a Lost or Forgotten Password	34-3
Procedure with Password Recovery Enabled	34-4
Procedure with Password Recovery Disabled	34-6
Preventing Autonegotiation Mismatches	34-7
SFP Module Security and Identification	34-8
Using Ping	34-8
Understanding Ping	34-8
Executing Ping	34-9
Using Layer 2 Traceroute	34-10
Understanding Layer 2 Traceroute	34-10
Usage Guidelines	34-10
Displaying the Physical Path	34-11
Using IP Traceroute	34-11
Understanding IP Traceroute	34-11
Executing IP Traceroute	34-12
Using Debug Commands	34-13
Enabling Debugging on a Specific Feature	34-13
Enabling All-System Diagnostics	34-14
Redirecting Debug and Error Message Output	34-14
Using the show platform forward Command	34-14
Using the crashinfo File	34-17

APPENDIX A

Supported MIBs A-1

MIB List	A-1
Using FTP to Access the MIB Files	A-3

APPENDIX B

Working with the Cisco IOS File System, Configuration Files, and Software Images B-1

- Working with the Flash File System B-1
 - Displaying Available File Systems B-2
 - Setting the Default File System B-3
 - Displaying Information about Files on a File System B-3
 - Changing Directories and Displaying the Working Directory B-3
 - Creating and Removing Directories B-4
 - Copying Files B-4
 - Deleting Files B-5
 - Creating, Displaying, and Extracting tar Files B-5
 - Creating a tar File B-6
 - Displaying the Contents of a tar File B-6
 - Extracting a tar File B-7
 - Displaying the Contents of a File B-7
- Working with Configuration Files B-8
 - Guidelines for Creating and Using Configuration Files B-8
 - Configuration File Types and Location B-9
 - Creating a Configuration File By Using a Text Editor B-9
 - Copying Configuration Files By Using TFTP B-10
 - Preparing to Download or Upload a Configuration File By Using TFTP B-10
 - Downloading the Configuration File By Using TFTP B-11
 - Uploading the Configuration File By Using TFTP B-11
 - Copying Configuration Files By Using FTP B-12
 - Preparing to Download or Upload a Configuration File By Using FTP B-12
 - Downloading a Configuration File By Using FTP B-13
 - Uploading a Configuration File By Using FTP B-14
 - Copying Configuration Files By Using RCP B-15
 - Preparing to Download or Upload a Configuration File By Using RCP B-16
 - Downloading a Configuration File By Using RCP B-16
 - Uploading a Configuration File By Using RCP B-17
 - Clearing Configuration Information B-18
 - Clearing the Startup Configuration File B-18
 - Deleting a Stored Configuration File B-19
- Working with Software Images B-19
 - Image Location on the Switch B-20
 - tar File Format of Images on a Server or Cisco.com B-20

Copying Image Files By Using TFTP	B-21
Preparing to Download or Upload an Image File By Using TFTP	B-21
Downloading an Image File By Using TFTP	B-22
Uploading an Image File By Using TFTP	B-23
Copying Image Files By Using FTP	B-24
Preparing to Download or Upload an Image File By Using FTP	B-24
Downloading an Image File By Using FTP	B-25
Uploading an Image File By Using FTP	B-27
Copying Image Files By Using RCP	B-28
Preparing to Download or Upload an Image File By Using RCP	B-28
Downloading an Image File By Using RCP	B-29
Uploading an Image File By Using RCP	B-31

APPENDIX C

Unsupported Commands in Cisco IOS Release 12.1(14)AX C-1

Access Control Lists	C-1
Unsupported Privileged EXEC Commands	C-1
Unsupported Global Configuration Commands	C-1
ARP	C-1
Unsupported Global Configuration Commands	C-1
Unsupported Interface Configuration Commands	C-2
Debug	C-2
Unsupported Privileged EXEC Commands	C-2
FallBack Bridging	C-2
Unsupported Privileged EXEC Commands	C-2
Unsupported Global Configuration Commands	C-2
Unsupported Interface Configuration Commands	C-3
HSRP	C-4
Unsupported Global Configuration Commands	C-4
Unsupported Interface Configuration Commands	C-4
IGMP Snooping	C-4
Unsupported Global Configuration Commands	C-4
Interfaces	C-4
Unsupported Privileged EXEC Commands	C-4
Unsupported Global Configuration Commands	C-4
Unsupported Interface Configuration Commands	C-5
IP Multicast Routing	C-5
Unsupported Privileged EXEC Commands	C-5
Unsupported Global Configuration Commands	C-6
Unsupported Interface Configuration Commands	C-6

IP Unicast Routing	C-6
Unsupported Privileged EXEC or User EXEC Commands	C-6
Unsupported Global Configuration Commands	C-7
Unsupported Interface Configuration Commands	C-7
Unsupported BGP Router Configuration Commands	C-8
Unsupported Route Map Commands	C-8
Layer 2 Protocol Tunneling	C-8
Unsupported Interface Configuration Commands	C-8
Miscellaneous	C-9
Unsupported Global Configuration Commands	C-9
MPLS	C-9
MSDP	C-9
Unsupported Privileged EXEC Commands	C-9
Unsupported Global Configuration Commands	C-9
RADIUS	C-9
Unsupported Global Configuration Commands	C-9
SNMP	C-10
Unsupported Global Configuration Commands	C-10
Spanning Tree	C-10
Unsupported Global Configuration Commands	C-10
Unsupported Interface Configuration Commands	C-10
Virtual Forwarding Infrastructure (VFI)	C-10
Unsupported Global Configuration Commands	C-10
Unsupported Privileged EXEC Commands	C-10
VLAN	C-10
Unsupported vlan-config Commands	C-10
Unsupported Privileged EXEC Commands	C-11
Unsupported User EXEC Commands	C-11
VTP	C-11
Unsupported Privileged EXEC Commands	C-11



Preface

Audience

This guide is for the networking professional managing the Catalyst 3750 Metro switch, hereafter referred to as the *switch*. Before using this guide, you should have experience working with the Cisco IOS software and be familiar with the concepts and terminology of Ethernet and local area networking.

Purpose

The Catalyst 3750 Metro switch supports a set of service-provider features, including multiprotocol label switching (MPLS), Ethernet over MPLS (EoMPLS), hierarchical QoS, hierarchical virtual private LAN switching (H-VPLS), 802.1Q tunneling, and VLAN mapping. It also supports Layer 2+ features such as access control lists (ACLs) and quality of service (QoS) and Layer 3 features such as IP unicast routing, IP multicast routing, and fallback bridging.

This guide provides procedures for using the commands that have been created or changed for use with the switch. It does not provide detailed information about these commands. For detailed information about these commands, refer to the *Catalyst 3750 Metro Switch Command Reference* for this release. 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 or how to install your switch. For more information, refer to the *Catalyst 3750 Metro Switch System Message Guide* for this release and to the *Catalyst 3750 Metro Switch Hardware Installation Guide*.

Conventions

This publication 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 ([]) mean optional elements.
- Braces ({ }) group required choices, and vertical bars (|) separate the alternative elements.
- Braces and vertical bars within square brackets ({ | }) mean 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 timesavers 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.



Timesaver

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

These documents provide complete information about the switch and are available from this Cisco.com site:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat3750m/index.htm>

**Note**

Before installing, configuring, or upgrading the switch, refer to these documents:

- For initial configuration information, refer to the “Using Express Setup” chapter or the “Configuring the Switch with the CLI-Based Setup Program” appendix in the hardware installation guide.
- For upgrade information, refer to the “Downloading Software” section in the release notes.

You can order printed copies of documents with a DOC-xxxxxx= number from the Cisco.com sites and from the telephone numbers listed in the “[Ordering Documentation](#)” section on page xxxvi.

- *Release Notes for the Catalyst 3750 Metro Switch* (not orderable but available on Cisco.com)
- *Catalyst 3750 Metro Switch Software Configuration Guide* (order number DOC-7815870=)
- *Catalyst 3750 Metro Switch Command Reference* (order number DOC-7815871=)
- *Catalyst 3750 Metro Switch System Message Guide* (order number DOC-7815872=)
- *Catalyst 3750 Metro Switch Hardware Installation Guide* (order number DOC-7815869=)
- *Cisco Small Form-Factor Pluggable Modules Installation Notes* (order number DOC-7815160=)

Obtaining Documentation

Cisco documentation and additional literature are available on Cisco.com. Cisco also provides several ways to obtain technical assistance and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

Cisco.com

You can access the most current Cisco documentation on the World Wide Web at this URL:

<http://www.cisco.com/univercd/home/home.htm>

You can access the Cisco website at this URL:

<http://www.cisco.com>

International Cisco websites can be accessed from this URL:

http://www.cisco.com/public/countries_languages.shtml

Ordering Documentation

You can find instructions for ordering documentation at this URL:

http://www.cisco.com/univercd/cc/td/doc/es_inpc/pdi.htm

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Ordering tool:

<http://www.cisco.com/en/US/partner/ordering/index.shtml>

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, USA.) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

You can submit e-mail comments about technical documentation to bug-doc@cisco.com.

You can submit comments by using the response card (if present) behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Customer Document Ordering
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

For all customers, partners, resellers, and distributors who hold valid Cisco service contracts, the Cisco Technical Assistance Center (TAC) provides 24-hour-a-day, award-winning technical support services, online and over the phone. Cisco.com features the Cisco TAC website as an online starting point for technical assistance. If you do not hold a valid Cisco service contract, please contact your reseller.

Cisco TAC Website

The Cisco TAC website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The Cisco TAC website is available 24 hours a day, 365 days a year. The Cisco TAC website is located at this URL:

<http://www.cisco.com/tac>

Accessing all the tools on the Cisco TAC website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a login ID or password, register at this URL:

<http://tools.cisco.com/RPF/register/register.do>

Opening a TAC Case

Using the online TAC Case Open Tool is the fastest way to open P3 and P4 cases. (P3 and P4 cases are those in which your network is minimally impaired or for which you require product information.) After you describe your situation, the TAC Case Open Tool automatically recommends resources for an immediate solution. If your issue is not resolved using the recommended resources, your case will be assigned to a Cisco TAC engineer. The online TAC Case Open Tool is located at this URL:

<http://www.cisco.com/tac/caseopen>

For P1 or P2 cases (P1 and P2 cases are those in which your production network is down or severely degraded) or if you do not have Internet access, contact Cisco TAC by telephone. Cisco TAC engineers are assigned immediately to P1 and P2 cases to help keep your business operations running smoothly.

To open a case by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411 (Australia: 1 800 805 227)

EMEA: +32 2 704 55 55

USA: 1 800 553-2447

For a complete listing of Cisco TAC contacts, go to this URL:

<http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

TAC Case Priority Definitions

To ensure that all cases are reported in a standard format, Cisco has established case priority definitions.

Priority 1 (P1)—Your network is “down” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Priority 2 (P2)—Operation of an existing network is severely degraded, or significant aspects of your business operation are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Priority 3 (P3)—Operational performance of your network is impaired, but most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Priority 4 (P4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- Cisco Marketplace provides a variety of Cisco books, reference guides, and logo merchandise. Go to this URL to visit the company store:

<http://www.cisco.com/go/marketplace/>

- The *Cisco Product Catalog* describes the networking products offered by Cisco Systems, as well as ordering and customer support services. Access the Cisco Product Catalog at this URL:

http://www.cisco.com/en/US/products/products_catalog_links_launch.html

- *Cisco Press* publishes a wide range of general networking, training and certification titles. Both new and experienced user will benefit from these publications. For current Cisco Press titles and other information, go to Cisco Press online at this URL:

<http://www.ciscopress.com>

- *Packet* magazine is the Cisco quarterly publication that provides the latest networking trends, technology breakthroughs, and Cisco products and solutions to help industry professionals get the most from their networking investment. Included are networking deployment and troubleshooting tips, configuration examples, customer case studies, tutorials and training, certification information, and links to numerous in-depth online resources. You can access Packet magazine at this URL:

<http://www.cisco.com/packet>

- *iQ Magazine* is the Cisco bimonthly publication that delivers the latest information about Internet business strategies for executives. You can access iQ Magazine at this URL:

<http://www.cisco.com/go/iqmagazine>

- *Internet Protocol Journal* is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:

<http://www.cisco.com/ipj>

- Training—Cisco offers world-class networking training. Current offerings in network training are listed at this URL:

<http://www.cisco.com/en/US/learning/index.html>



Overview

This chapter provides this information about Catalyst 3750 Metro switch software:

- [Features, page 1-1](#)
- [Default Settings After Initial Switch Configuration, page 1-8](#)
- [Network Configuration Examples, page 1-11](#)
- [Where to Go Next, page 1-16](#)

Features



Note

Some features noted in this chapter are available only on the cryptographic (that is, supports encryption) versions of the switch software image. You must obtain authorization to use this feature and to download the cryptographic version of the software from Cisco.com. For more information, refer to the release notes for this release.

Catalyst 3750 Metro switches have these features:

- [Performance Features, page 1-2](#)
- [Management Options, page 1-2](#)
- [Manageability Features, page 1-3](#) (includes a feature requiring the cryptographic version of the switch software image)
- [Availability Features, page 1-3](#)
- [VLAN Features, page 1-4](#)
- [Layer 2 Virtual Private Network \(VPN\) Services, page 1-4](#)
- [Layer 3 VPN Services, page 1-5](#)
- [Security Features, page 1-5](#) (includes a feature requiring the cryptographic version of the switch software image)
- [QoS Features, page 1-6](#)
- [Layer 3 Features, page 1-7](#)
- [Monitoring Features, page 1-8](#)

Performance Features

- Autosensing of port speed and autonegotiation of duplex mode on all switch ports for optimizing bandwidth
- Automatic-medium-dependent interface crossover (Auto-MDIX) capability on 10/100 Mbps interfaces and 10/100/1000 BASE-T/TX small form-factor pluggable (SFP) interfaces that enables the interface to automatically detect the required cable connection type (straight through or crossover) and configure the connection appropriately
- IEEE 802.3X flow control on all ports (the switch does not send pause frames)
- EtherChannel for enhanced fault tolerance and for providing up to 2 Gbps (Gigabit EtherChannel) or 800 Mbps (Fast EtherChannel) full duplex of bandwidth between switches, routers, and servers
- Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP) for automatic creation of EtherChannel links
- Per-port storm control for preventing broadcast, multicast, and unicast storms
- Port blocking on forwarding unknown Layer 2 unknown unicast, multicast, and bridged broadcast traffic
- Cisco Group Management Protocol (CGMP) server support and Internet Group Management Protocol (IGMP) snooping for IGMP versions 1 and 2:
 - (For CGMP devices) CGMP for limiting multicast traffic to specified end stations and reducing overall network traffic
 - (For IGMP devices) IGMP snooping for efficiently forwarding multimedia and 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
- Switch Database Management (SDM) templates for allocating system resources to maximize support for user-selected features

Management Options

- CLI—The Cisco IOS command-line interface (CLI) software is enhanced to support the Catalyst 3750 Metro switch features. You can access the CLI either by connecting your management station directly to the switch console port or by using Telnet from a remote management station. For more information about the CLI, see [Chapter 2, “Using the Command-Line Interface.”](#)
- IE2100—Cisco Intelligence Engine 2100 Series Configuration Registrar is a network management device that works with embedded Cisco Networking Services (CNS) Agents in the switch software. You can automate initial configurations and configuration updates by generating switch-specific configuration changes, sending them to the switch, executing the configuration change, and logging the results.

For more information about IE2100, see [Chapter 4, “Configuring IE2100 CNS Agents.”](#)

- SNMP—Simple Network Management Protocol (SNMP) management applications such as CiscoWorks2000 LAN Management Suite (LMS) and HP OpenView. You can manage 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 remote monitoring (RMON) groups. For more information about using SNMP, see [Chapter 24, “Configuring SNMP.”](#)

Manageability Features

**Note**

The encrypted Secure Shell (SSH) feature listed in this section is available only on the cryptographic version of the switch software image.

- Cisco IE2100 Series CNS embedded agents for automating switch management, configuration storage, and delivery
- Dynamic Host Configuration Protocol (DHCP) for automating configuration of switch information (such as IP address, default gateway, host name, and Domain Name System [DNS] and TFTP server names)
- Directed unicast requests to a DNS server for identifying a switch through its IP address and its corresponding host name and to a TFTP server for administering software upgrades from a TFTP server
- 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
- Cisco IOS File System (IFS) for providing a single interface to all file systems that the switch uses
- In-band management access for up to 16 simultaneous Telnet connections for multiple CLI-based sessions over the network
- In-band management access for up to five simultaneous, encrypted SSH connections for multiple CLI-based sessions over the network (requires the cryptographic version of the switch software image)
- In-band management access through SNMP versions 1 and 2c, and 3 get and set requests
- Out-of-band management access through the switch console port to a directly attached terminal or to a remote terminal through a serial connection or a modem

Availability Features

- HSRP for command switch and Layer 3 router redundancy
- UniDirectional Link Detection (UDLD) and aggressive UDLD for detecting and disabling unidirectional links on fiber-optic interfaces caused by incorrect fiber-optic wiring or port faults
- IEEE 802.1D Spanning Tree Protocol (STP) for redundant backbone connections and loop-free networks. STP has these features:
 - Up to 128 spanning-tree instances supported
 - Per-VLAN spanning-tree plus (PVST+) for balancing load across VLANs
 - Rapid PVST+ for balancing load across VLANs and providing rapid convergence of spanning-tree instances
 - 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 Spanning Tree Protocol (MSTP) for grouping VLANs into a spanning-tree instance and for providing multiple forwarding paths for data traffic and load balancing and IEEE 802.1w Rapid Spanning Tree Protocol (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 PVST+, rapid-PVST+, and MSTP mode:
 - 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 bridge protocol data units (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
- Equal-cost routing for link-level and switch-level redundancy

VLAN Features

- Support for up to 1005 VLANs for assigning users to VLANs associated with appropriate network resources, traffic patterns, and bandwidth
- Support for VLAN IDs in the full 1 to 4094 range allowed by the IEEE 802.1Q standard
- VLAN Query Protocol (VQP) for dynamic VLAN membership
- Inter-Switch Link (ISL) and IEEE 802.1Q trunking encapsulation 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
- Dynamic Trunking Protocol (DTP) for negotiating trunking on a link between two devices and for negotiating the type of trunking encapsulation (802.1Q or ISL) to be used
- VLAN Trunking Protocol (VTP) and VTP pruning for reducing network traffic by restricting flooded traffic to links destined for stations receiving the traffic
- Voice VLAN for creating subnets for voice traffic from Cisco IP Phones
- VLAN 1 minimization for reducing 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 on the trunk, but the switch CPU continues to send and receive control protocol frames.
- VLAN mapping on enhanced-services (ES) ports to translate customer VLANs to service-provider VLANs for transporting packets across the service-provider network without affecting the customer VLAN IDs
- Custom ethertype to enable the user to change the ethertype value on a port to any value to direct the tagged and untagged traffic to different VLANs (ES ports only)

Layer 2 Virtual Private Network (VPN) Services

- 802.1Q tunneling so that customers with users at remote sites across a service-provider network can keep VLANs segregated from other customers, and Layer 2 protocol tunneling to ensure that the customer's network has complete STP, CDP, and VTP information about all users
- Layer 2 protocol tunneling bypass feature to provide interoperability with third-party vendors
- Intelligent 802.1Q tunneling QoS, the ability to copy the inner cost-of-service (CoS) value to the outer CoS value for 802.1Q tunneling
- Support for Ethernet over multiprotocol layer switching (EoMPLS) tunneling mechanism for transporting Ethernet frames over a service-provider MPLS network

Layer 3 VPN Services

- Support for MPLS VPNs provides the capability to deploy and administer scalable Layer 3 VPN services to business customers. Each VPN is associated with one or more VPN routing/forwarding (VRF) instances that include routing and forwarding tables and rules that define the VPN membership. (MPLS VPNs are supported only on ES ports.)
- Multiple VPN multi-VRF instances in customer edge (CE) devices to allow service providers to support multiple VPNs and to overlap IP addresses between VPNs.

Security Features



Note

The Kerberos feature listed in this section is available only on the cryptographic version of the switch software image.

- Password-protected access (read-only and read-write access) to management interfaces for protection against unauthorized configuration changes
- Multilevel security for a choice of security level, notification, and resulting actions
- Static MAC addressing for ensuring security
- Protected port option for restricting the forwarding of traffic to designated ports on the same switch
- 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
- BPDU guard for shutting down a Port Fast-configured port when an invalid configuration occurs
- Standard and extended IP access control lists (ACLs) for defining security policies in both directions on routed interfaces (router ACLs) and VLANs and inbound on Layer 2 interfaces (port ACLs)
- Extended MAC access control lists for defining security policies in the inbound direction on Layer 2 interfaces
- VLAN ACLs (VLAN maps) for providing intra-VLAN security by filtering traffic based on information in the MAC, IP, and TCP/User Datagram Protocol (UDP) headers
- Source and destination MAC-based ACLs for filtering non-IP traffic
- IEEE 802.1x port-based authentication to prevent unauthorized devices (clients) from gaining access to the network
- TACACS+, a proprietary feature for managing network security through a TACACS server
- RADIUS for verifying the identity of, granting access to, and tracking the actions of remote users through authentication, authorization, and accounting (AAA) services
- Kerberos security system to authenticate requests for network resources by using a trusted third party (requires the cryptographic [that is, supports encryption] version of the switch software image)
- Password recovery disable capability to protect access to switches at customer sites

QoS Features

- Standard QoS to classify, police, mark, queue, and schedule incoming traffic on a standard port or on an ES port, as well as queue and schedule outgoing traffic on a standard port

Classification

- IP type-of-service/Differentiated Services Code Point (IP ToS/DSCP) and 802.1p CoS marking priorities on a per-port basis for protecting the performance of mission-critical applications
- IP ToS/DSCP and 802.1p CoS marking based on flow-based packet classification (classification based on information in the MAC, IP, and TCP/UDP headers) for high-performance quality of service at the network edge, allowing for differentiated service levels for different types of network traffic and for prioritizing mission-critical traffic in the network
- Trusted port states (CoS, DSCP, and IP precedence) within a QoS domain and with a port bordering another QoS domain
- Trusted boundary for detecting the presence of a Cisco IP phone, trusting the CoS value received, and ensuring port security

Policing and out-of-profile marking

- Traffic-policing policies on the switch port for managing how much of the port bandwidth should be allocated to a specific traffic flow (single-rate traffic policing)
- Aggregate policing for policing traffic flows in aggregate to restrict specific applications or traffic flows to metered, predefined rates
- Out-of-profile markdown for packets that exceed bandwidth utilization limits (drop policy actions are passing through the packet without modification, marking down the assigned DSCP in the packet, or dropping the packet)

Ingress queueing and scheduling

- Two configurable ingress queues for user traffic (one queue can be the priority queue)
- Weighted tail drop (WTD) as the congestion-avoidance mechanism for managing the queue lengths and providing drop precedences for different traffic classifications
- Shaped round robin (SRR) as the scheduling service for determining the rate at which packets are dequeued to the internal ring (sharing is the only supported mode on ingress queues)

Egress queues and scheduling

- Four egress queues per port
 - WTD as the congestion-avoidance mechanism for managing the queue lengths and providing drop precedences for different traffic classifications
 - SRR as the scheduling service for determining the rate at which packets are dequeued to the egress interface (shaping or sharing is supported on egress queues). Shaped egress queues are guaranteed but limited to using a share of port bandwidth. Shared egress queues are also guaranteed a configured share of bandwidth, but can use more than the guarantee if other queues become empty and do not use their share of the bandwidth.
- Hierarchical QoS on ES ports to classify, police, mark, queue, and schedule outgoing traffic

Classification

- Three QoS configuration levels in the hierarchy: class, VLAN, and physical interface
- Classification based on the CoS value, the DSCP value, the IP precedence value, the MPLS experimental (EXP) bits, or the VLAN

Policing and out-of-profile marking

- Two-rate traffic policing based on the committed information rate (CIR) and the peak information rate (PIR)
- Out-of-profile markdown for packets that exceed bandwidth utilization limits (policy actions are to send packets that conform without modification, to mark down the priority of packets that exceed, and to drop packets that violate)

Egress queueing and scheduling

- Each packet assigned to an egress queue based on the traffic class and VLAN
 - Weighted Random Early Detection (WRED) as the congestion-avoidance mechanism
 - Class-based weighted fair queueing (CBWFQ) as a queue scheduling management feature to provide guaranteed bandwidth to particular traffic classes, such as voice, that are delay sensitive, while still fairly serving all other traffic in the network
 - Low-latency queueing (LLQ) as a scheduling congestion-management feature to provide strict-priority queueing for a traffic class and to enable delay-sensitive data, such as voice, to be sent before packets in other queues are sent
 - Traffic shaping to decrease the burstiness of Internet traffic
- Automatic QoS (auto-QoS) to simplify the deployment of existing QoS features by classifying traffic and configuring egress queues (voice over IP only)

Layer 3 Features

- HSRP for Layer 3 router redundancy
- IP routing protocols for load balancing and for constructing scalable, routed backbones:
 - RIP versions 1 and 2
 - OSPF
 - Interior Gateway Routing Protocol (IGRP) and Enhanced IGRP (EIGRP)
 - Border Gateway Protocol (BGP) Version 4
 - International Organization for Standardization (ISO) Connectionless Network Service (CLNS). The switch supports ISO Interior Gateway Routing Protocol (IGRP) and Intermediate System-to-Intermediate System (IS-IS) routing
- IP routing between VLANs (inter-VLAN routing) for full Layer 3 routing between two or more VLANs, allowing each VLAN to maintain its own autonomous data-link domain
- Policy-based routing (PBR) for configuring defined policies for traffic flows
- Fallback bridging for forwarding non-IP traffic between two or more VLANs
- Static IP routing for manually building a routing table of network path information
- Equal-cost routing for load balancing and redundancy
- Internet Control Message Protocol (ICMP) and ICMP Router Discovery Protocol (IRDP) for using router advertisement and router solicitation messages to discover the addresses of routers on directly attached subnets
- Protocol-Independent Multicast (PIM) for multicast routing within the network, allowing for devices in the network to receive the multicast feed requested and for switches not participating in the multicast to be pruned. Includes support for PIM sparse mode (PIM-SM), PIM dense mode (PIM-DM), and PIM sparse-dense mode.
- Multicast Source Discovery Protocol (MSDP) for connecting multiple PIM-SM domains

- Distance Vector Multicast Routing Protocol (DVMRP) tunnelling for interconnecting two multicast-enabled networks across non-multicast networks
- DHCP relay for forwarding UDP broadcasts, including IP address requests, from DHCP clients

Monitoring Features

- Switch LEDs that provide port- and switch-level status
- MAC address notification traps and RADIUS accounting for tracking users on a network by storing the MAC addresses that the switch has learned or removed
- Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) for traffic monitoring on any standard port or VLAN.



Note An ES port cannot be a SPAN source.

- SPAN and RSPAN support of Intrusion Detection Systems (IDS) to monitor, repel, and report network security violations
- Four groups (history, statistics, alarms, and events) of embedded RMON agents for network monitoring and traffic analysis
- 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

Default Settings After Initial Switch Configuration

The switch is designed for plug-and-play operation, requiring only that you 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 change the interface-specific and system-wide settings.

If you do not configure the switch at all, the switch operates with the default settings listed in [Table 1-1](#). This table lists the key software features, their defaults, and where to find more information about the features.

For information about setting up the initial switch configuration (by using Express Setup or the CLI setup program) and assigning basic IP information to the switch, refer to the hardware installation guide.

Table 1-1 Default Settings After Initial Switch Configuration

Feature	Default Setting	More information in...
Switch IP address, subnet mask, and default gateway	0.0.0.0	Chapter 3, “Assigning the Switch IP Address and Default Gateway”
Domain name	None	
DHCP	DHCP client enabled	

Table 1-1 *Default Settings After Initial Switch Configuration (continued)*

Feature	Default Setting	More information in...
Passwords	None defined	Chapter 5, “Administering the Switch”
TACACS+	Disabled	
RADIUS	Disabled	
System name and prompt	<i>Switch</i>	
NTP	Enabled	
DNS	Enabled	
802.1x	Disabled	Chapter 8, “Configuring 802.1x Port-Based Authentication”
Port parameters		
Operating mode	Layer 2 (switchport)	Chapter 9, “Configuring Interface Characteristics”
Port enable state	All ports are enabled.	
Interface speed and duplex mode	Autonegotiate	
Auto MDIX	Disabled	
Flow control	Off	
VLANs		
Default VLAN	VLAN 1	Chapter 10, “Configuring VLANs”
VLAN trunking	Dynamic auto (DTP)	
Trunk encapsulation	Negotiate	
VTP mode	Server	Chapter 11, “Configuring VTP”
VTP version	1	
Voice VLAN	Disabled	Chapter 12, “Configuring Voice VLAN”
Tunneling		
802.1Q tunneling	Disabled	Chapter 13, “Configuring IEEE 802.1Q and Layer 2 Protocol Tunneling”
VLAN mapping	Disabled	
Layer 2 protocol tunneling	Disabled	
Spanning Tree Protocol		
STP	PVST+ enabled on VLAN 1	Chapter 14, “Configuring STP”
MSTP	Disabled	Chapter 15, “Configuring MSTP”
Optional spanning-tree features	Disabled	Chapter 16, “Configuring Optional Spanning-Tree Features”
IGMP snooping		
IGMP snooping	Enabled	Chapter 17, “Configuring IGMP Snooping and MVR”
IGMP filters	None applied	
MVR	Disabled	

Table 1-1 *Default Settings After Initial Switch Configuration (continued)*

Feature	Default Setting	More information in...
Port-based Traffic		
Broadcast, multicast, and unicast storm control	Disabled	Chapter 18, “Configuring Port-Based Traffic Control”
Protected ports	None defined	
Unicast and multicast traffic flooding	Not blocked	
Secure ports	None configured	
CDP	Enabled	Chapter 19, “Configuring CDP”
UDLD	Disabled	Chapter 20, “Configuring UDLD”
SPAN and RSPAN	Disabled	Chapter 21, “Configuring SPAN and RSPAN”
RMON	Disabled	Chapter 22, “Configuring RMON”
Syslog messages	Enabled; displayed on the console	Chapter 23, “Configuring System Message Logging”
SNMP	Enabled; version 1	Chapter 24, “Configuring SNMP”
ACLs	None configured	Chapter 25, “Configuring Network Security with ACLs”
QoS	Disabled	Chapter 26, “Configuring QoS”
EtherChannels	None configured	Chapter 27, “Configuring EtherChannels”
IP unicast routing		
IP routing (and routing protocols)	Disabled	Chapter 28, “Configuring IP Unicast Routing”
Multi-VRF-CE	Disabled	
MPLS services		
Label switching	Globally enabled; disabled per interface	Chapter 30, “Configuring MPLS and EoMPLS”
EoMPLS	Not configured	
MPLS QoS	Disabled	
HSRP groups	None configured	Chapter 29, “Configuring HSRP”
IP multicast routing	Disabled on all interfaces	Chapter 31, “Configuring IP Multicast Routing”
MSDP	Disabled	Chapter 32, “Configuring MSDP”
Fallback bridging	Not configured	Chapter 33, “Configuring Fallback Bridging”

Network Configuration Examples

This section provides network configuration concepts and includes examples of using the switch to create dedicated network segments and interconnecting the segments through Fast Ethernet and Gigabit Ethernet connections.

- [“Multidwelling or Ethernet-to-the Subscriber Network” section on page 1-11](#)
- [“Ethernet Broadband Aggregation Network” section on page 1-13](#)
- [“Layer 2 VPN Application” section on page 1-14](#)
- [“Layer 3 VPN Application” section on page 1-15](#)

Multidwelling or Ethernet-to-the Subscriber Network

Figure 1-1 shows a Gigabit Ethernet ring for a residential location serving multitenant units using Catalyst 3750 Metro switches connected through 1000BASE-X SFP module ports. Catalyst 3750 Metro switches used as residential switches provide customers with high-speed connections to the service provider point-of presence (POP). Catalyst 2950 Long-Reach Ethernet (LRE) switches also can be used as residential switches for customers requiring connectivity through existing phone lines. The Catalyst 2950 LRE switches can then connect to another residential switch, such as a Catalyst 3750 Metro switch. For more information about the Catalyst LRE switches and LRE information, refer to the Catalyst 2950 LRE documentation set.

All ports on the residential switches (and Catalyst 2950 LRE switches if they are included) are configured as 802.1Q trunks with Private VLAN Edge (protected port) and STP root guard features enabled. The protected-port feature provides security and isolation between ports on the switch, ensuring that subscribers cannot view packets destined for other subscribers. STP root guard prevents unauthorized devices from becoming the STP root switch. All ports have IGMP snooping or CGMP enabled for multicast traffic management. ACLs on the uplink ports to the aggregating switches provide security and bandwidth management.


The aggregating switches and routers have HSRP enabled for load balancing and redundant connectivity to guarantee mission-critical traffic. This ensures connectivity to the Internet, WAN, and mission-critical network resources in case one of the routers or switches fails.

When an end station in one VLAN needs to communicate with an end station in another VLAN, a router or switch routes the traffic to the appropriate destination VLAN, providing inter-VLAN routing. VLAN access control lists (VLAN maps) provide intra-VLAN security and prevent unauthorized users from accessing critical pieces of the network.

In addition to inter-VLAN routing, the switch QoS mechanisms such as DSCP prioritize the different types of network traffic to deliver high-priority traffic in a predictable manner. If congestion occurs, QoS drops low-priority traffic to allow delivery of high-priority traffic.

The routers also provide firewall services, Network Address Translation (NAT) services, voice-over-IP (VoIP) gateway services, and WAN and Internet access.

Figure 1-1 Catalyst 3750 Metro Switches in a Multidwelling Configuration




Ethernet Broadband Aggregation Network

Digital subscriber line (DSL) and fiber-to-the-house (FTTH) services provide high performance and increased bandwidth to the home or small-business customer for use in IP telephones, televisions, or PCs.

Catalyst 3750 Metro switch hierarchical QoS features allow service providers to support differentiated services with different levels of services for multiple customers. The configuration is applicable using DSL, sending packets through a digital subscriber line access multiplexer (DSLAM) and DSL modem to the residence, or using fiber-optic lines through a media converter to the residence. The end-user device can be a television, PC, or IP telephone.

Figure 1-2 Broadband Aggregation Network Using DSL or FTTH



Layer 2 VPN Application


You can use Catalyst 3750 Metro switches to form Layer 2 VPNs so that customers at different locations can exchange information through a service-provider network, without requiring dedicated connections. IEEE 802.1Q and Layer 2 protocol tunneling are features designed for service providers who carry traffic of multiple customers across their networks and are required to maintain the VLAN and Layer 2 protocol configurations of each customer without impacting the traffic of other customers.

The Catalyst 3750 Metro switches are used as the provider edge (PE) switches at the both edges of the provider network connected to customer premises equipment (CPE) switches. The PE switches tag packets entering the service-provider network with the customer VLAN ID. VLAN mapping translates each customer VLAN ID to a service-provider VLAN ID for transport across the service-provider network. At the egress PE interface, the egress PE switch restores the original VLAN ID numbers for the customer's network.

The service provider can use 802.1Q tunneling or EoMPLS to provide Layer 2 VPN services. If the service-provider network is an MPLS cloud, and EoMPLS is configured as the point-to-point protocol, MPLS tags are added at the ingress PE ES port that connects to the MPLS network. The MPLS tags are removed at the ES port of the remote PE device.

See [Chapter 13, “Configuring IEEE 802.1Q and Layer 2 Protocol Tunneling,”](#) and [Chapter 30, “Configuring MPLS and EoMPLS,”](#) for more information on configuring these features.

Figure 1-3 Layer 2 VPN Configuration



Layer 3 VPN Application

Layer 3 VPN services can use multi-VRF-CE or MPLS VPNs to deploy and administer scalable Layer 3 VPN services to business customers. A Layer 3 VPN is a secure IP-based network that shares resources on one or more physical networks. It contains geographically dispersed sites that can communicate securely over a shared backbone.


Figure 1-4 illustrates a typical MPLS VPN configuration. The CE devices (which can be Catalyst 3750 Metro switches or other Layer 3 switches) use a routing protocol, such as RIP, EIGRP, OSPF, IS-IS, or static routing, to forward packets from customer VPNs to the Catalyst 3750 Metro PE devices at the edge of the MPLS network. The PE device is configured with multiprotocol BGP (MP-BGP), and a route distinguisher that is associated with the customer's VPN. The PE device converts this information to a VPN-IPv4 format and adds layer distribution protocol (LDP) labels to establish VPN routes.

VPN routes are distributed over the MPLS network using MP-BGP, which also distributes the labels associated with each VPN route. MPLS VPN depends on VPN routing and forwarding (VRF) support to isolate the routing domains from each other.

When an MPLS-VPN packet is received on a port, the CE switch looks up the labels in the routing table to determine what to do with the packet. A PE router binds a label to each customer prefix learned from a CE device and includes the label in the prefix that it advertises to other PE routers. When a PE router forwards a packet across the provider network, it labels the packet with the label learned from the destination PE router. When the destination PE router receives the labeled packet, it examines the label and uses it to direct the packet to the correct CE device.

Only the PE routers at each end of the MPLS network maintain the VPN routes for VPN members. Provider routers in the core network do not maintain the VPN routes. This ensures the security of customer VPNs and isolates them from other customer packets that are carried across the service-provider MPLS network.

Figure 1-4 MPLS VPN Configuration



101097

See [Chapter 30, “Configuring MPLS and EoMPLS,”](#) for more information on configuring MPLS VPN.

Where to Go Next

Before configuring the switch, review these sections for startup information:

- [Chapter 2, “Using the Command-Line Interface”](#)
- [Chapter 3, “Assigning the Switch IP Address and Default Gateway”](#)



Using the Command-Line Interface

This chapter describes the Cisco IOS command-line interface (CLI) and how to use it to configure your Catalyst 3750 Metro switch. It contains these sections:

- [Understanding Command Modes, page 2-1](#)
- [Understanding the Help System, page 2-3](#)
- [Understanding Abbreviated Commands, page 2-4](#)
- [Understanding no and default Forms of Commands, page 2-4](#)
- [Understanding CLI Error Messages, page 2-5](#)
- [Using Command History, page 2-5](#)
- [Using Editing Features, page 2-6](#)
- [Searching and Filtering Output of show and more Commands, page 2-9](#)
- [Accessing the CLI, page 2-9](#)

Understanding 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 2-1](#) 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 2-1 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 • 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.

Table 2-1 Command Mode Summary (continued)

Mode	Access Method	Prompt	Exit Method	About This Mode
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 Ethernet ports. For information about defining interfaces, see the “Using Interface Configuration Mode” section on page 9-6. To configure multiple interfaces with the same parameters, see the “Configuring a Range of Interfaces” section on page 9-8.
Line configuration	While in global configuration mode, specify a line with the line vty or line console 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.

Understanding the Help System

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 2-2](#).

Table 2-2 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</i> <Tab>	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> ?

Table 2-2 *Help Summary (continued)*

Command	Purpose
<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

Understanding Abbreviated Commands

You need 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 in an abbreviated form:

```
Switch# show conf
```

Understanding 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 Error Messages

Table 2-3 lists some error messages that you might encounter while using the CLI to configure your switch.

Table 2-3 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.
% 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 software provides a history or record of commands that you have entered. The command history feature is particularly useful for recalling long or complex commands or entries, including access lists. You can customize this feature to suit your needs as described in these sections:

- [Changing the Command History Buffer Size, page 2-5](#) (optional)
- [Recalling Commands, page 2-6](#) (optional)
- [Disabling the Command History Feature, page 2-6](#) (optional)

Changing the Command History Buffer Size

By default, the switch records ten command lines in its history buffer. You can alter this number for a current terminal session or for all sessions on a particular line. These procedures are optional.

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 2-4](#). These actions are optional.

Table 2-4 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 controlled by 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. You can disable it for the current terminal session or for the command line. These procedures are optional.

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, page 2-7](#) (optional)
- [Editing Commands through Keystrokes, page 2-7](#) (optional)
- [Editing Command Lines that Wrap, page 2-8](#) (optional)

Enabling and Disabling Editing Features

Although enhanced editing mode is automatically enabled, you can disable it, re-enable it, or configure a specific line to have enhanced editing. These procedures are optional.

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

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

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
```

Editing Commands through Keystrokes

Table 2-5 shows the keystrokes that you need to edit command lines. These keystrokes are optional.

Table 2-5 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.

Table 2-5 Editing Commands through Keystrokes (continued)

Capability	Keystroke ¹	Purpose
	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 .	
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. The keystroke actions are optional.

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 2-7](#).

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. Using these commands is optional.

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
GigabitEthernet1/0/1 is up, line protocol is down
GigabitEthernet1/0/2 is up, line protocol is up
```

Accessing the CLI

You can access the CLI through a console connection, through Telnet, or by using the browser.

Accessing the CLI through a Console Connection or through Telnet

Before you can access the CLI, you must connect a terminal or PC to the switch console port and power on the switch as described in the hardware installation guide that shipped with your switch. Then, to understand the boot process and the options available for assigning IP information, see [Chapter 3, “Assigning the Switch IP Address and Default Gateway.”](#)

If your switch is already configured, you can access the CLI through a local console connection or through a remote Telnet session, but your switch must first be configured for this type of access. For more information, see the [“Setting a Telnet Password for a Terminal Line” section on page 7-6.](#)

You can use one of these methods to establish a connection with the switch:

- Connect the switch console port to a management station or dial-up modem. For information about connecting to the console port, refer to the switch hardware installation guide.
- Use 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 [“Setting a Telnet Password for a Terminal Line” section on page 7-6.](#) 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 7-37.](#) The switch supports up to five simultaneous secure SSH sessions.

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

Accessing the CLI from a Browser

Before performing this procedure, make sure that you have assigned IP information as described in the switch hardware installation guide. You also must assign a Telnet password to the switch as described in [“Setting a Telnet Password for a Terminal Line” section on page 7-6](#)

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

-
- Step 1** In the **URL** field, enter the IP address of the switch.
 - Step 2** When the Cisco Systems Access page appears, click **Telnet** to start a Telnet session.
 - Step 3** Enter the switch password.

The user EXEC prompt appears on the management station.



Assigning the Switch IP Address and Default Gateway

This chapter describes how to create the initial switch configuration (for example, assigning the switch IP address and default gateway information) for the Catalyst 3750 Metro switch by using a variety of automatic and manual methods. It also describes how to modify the switch startup configuration.



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, page 3-1](#)
- [Assigning Switch Information, page 3-2](#)
- [Checking and Saving the Running Configuration, page 3-10](#)
- [Modifying the Startup Configuration, page 3-10](#)
- [Scheduling a Reload of the Software Image, page 3-15](#)

Understanding the Boot Process

To start your switch, you need to follow the procedures in the hardware installation guide about installing and powering on the switch, and setting up the initial configuration (IP address, subnet mask, default gateway, secret and Telnet passwords, and so forth) of the 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. For more information, see the [“Recovering from Corrupted Software By Using the XMODEM Protocol” section on page 34-2](#) and the [“Recovering from a Lost or Forgotten Password” section on page 34-3](#).

**Note**

You can disable password recovery. For more information, see the [“Disabling Password Recovery” section on page 7-5](#).

Before you can assign switch information using the CLI or a DHCP server, make sure you have connected a PC or terminal to the console port, and configured the PC or terminal-emulation software baud rate and character format to match these of the switch console port:

- Baud rate default is 9600.
- Data bits default is 8.
- Stop bits default is 1.
- Parity settings default is none.

**Note**

If you are using Express Setup, do not connect any devices to the switch before starting Express Setup.

Assigning Switch Information

You can assign IP information using Express Setup, using the switch CLI-based setup program, through a DHCP server, or manually.

Use the Express Setup or the CLI-based setup program if you are a new user and want to be prompted for specific IP information. With these programs, you can also configure a default gateway, a host name, a switch password, and other information. You also have the option of assigning a Telnet password (to provide security during remote management). For more information about the setup programs, refer to the hardware installation guide.

Use a DHCP server for centralized control and automatic assignment of IP information once the server is configured.

**Note**

If you are using DHCP, do not respond to any of the questions in the setup program until the switch receives the dynamically-assigned IP address and reads the configuration file.

Use the manual method of configuration if you are an experienced user familiar with the switch configuration steps; otherwise, use the setup program described earlier.

This section contains this configuration information:

- [Default Switch Information, page 3-3](#)
- [Understanding DHCP-Based Autoconfiguration, page 3-3](#)
- [Configuring DHCP-Based Autoconfiguration, page 3-4](#)
- [Manually Assigning IP Information, page 3-9](#)

Default Switch Information

[Table 3-1](#) shows the default switch information.

Table 3-1 *Default Switch Information*

Feature	Default Setting
IP address and subnet mask	No IP address or subnet mask are defined.
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 password	No password is defined.

Understanding DHCP-Based Autoconfiguration

The DHCP provides configuration information to Internet hosts and internetworking devices. This protocol consists of two components: one for delivering configuration parameters from a DHCP server to a device and a mechanism for allocating network addresses to devices. DHCP is built on a client-server model, in which designated DHCP servers allocate network addresses and deliver configuration parameters to dynamically configured devices.

During DHCP-based autoconfiguration, your switch (DHCP client) is automatically configured at startup with IP address information and a configuration file.

With DHCP-based autoconfiguration, no DHCP client-side configuration is needed on your switch. However, you need to configure the DHCP server for various lease options associated with IP addresses. If you are using DHCP to relay the configuration file location on the network, you might also need to configure a TFTP server and a Domain Name System (DNS) server.

The DHCP server can be on the same LAN or on a different LAN than the switch. If the DHCP server is running on a different LAN, you should configure a DHCP relay. A relay device forwards broadcast traffic between two directly connected LANs. A router does not forward broadcast packets, but it forwards packets based on the destination IP address in the received packet.


DHCP-based autoconfiguration replaces the BOOTP client functionality on your switch.

DHCP Client Request Process

When you boot your switch, the DHCP client is invoked and automatically requests configuration information from a DHCP server when the configuration file is not present on the switch.

[Figure 3-1](#) shows the sequence of messages that are exchanged between the DHCP client and the DHCP server.

Figure 3-1 DHCP Client and Server Message Exchange



The client, Switch A, broadcasts a DHCPDISCOVER message to locate a DHCP server. The DHCP server offers configuration parameters (such as an IP address, subnet mask, gateway IP address, DNS IP address, a lease for the IP address, and so forth) to the client in a DHCPOFFER unicast message.

In a DHCPREQUEST broadcast message, the client returns a formal request for the offered configuration information to the DHCP server. The formal request is broadcast so that all other DHCP servers that received the DHCPDISCOVER broadcast message from the client can reclaim the IP addresses that they offered to the client.

The DHCP server confirms that the IP address has been allocated to the client by returning a DHCPACK unicast message to the client. With this message, the client and server are bound, and the client uses configuration information received from the server. The amount of information the switch receives depends on how you configure the DHCP server. For more information, see the [“Configuring the DHCP Server”](#) section on page 3-5.

If the configuration parameters sent to the client in the DHCPOFFER unicast message are invalid (a configuration error exists), the client returns a DHCPDECLINE broadcast message to the DHCP server.

The DHCP server sends the client a DHCPNAK denial broadcast message, which means that the offered configuration parameters have not been assigned, that an error has occurred during the negotiation of the parameters, or that the client has been slow in responding to the DHCPOFFER message (the DHCP server assigned the parameters to another client).

A DHCP client might receive offers from multiple DHCP or BOOTP servers and can accept any of the offers; however, the client usually accepts the first offer it receives. The offer from the DHCP server is not a guarantee that the IP address is allocated to the client; however, the server usually reserves the address until the client has had a chance to formally request the address. If the switch accepts replies from a BOOTP server and configures itself, the switch broadcasts, instead of unicasts, TFTP requests to obtain the switch configuration file.

Configuring DHCP-Based Autoconfiguration

These sections describe how to configure DHCP-based autoconfiguration:

- [Configuring the DHCP Server, page 3-5](#)
- [Configuring the TFTP Server, page 3-5](#)
- [Configuring the DNS, page 3-6](#)

- [Configuring the Relay Device, page 3-6](#)
- [Configuring the Relay Device, page 3-6](#)
- [Obtaining Configuration Files, page 3-7](#)
- [Example Configuration, page 3-8](#)

If your DHCP server is a Cisco device, or if you are configuring the switch as a DHCP server, refer to the “*IP Addressing and Services*” section in the *Cisco IOS IP and IP Routing Configuration Guide for Cisco IOS Release 12.1* for additional information about configuring DHCP.

Configuring the DHCP Server

You should configure the DHCP server with reserved leases that are bound to each switch by the switch hardware address.

If you want the switch to receive IP address information, you must configure the DHCP server with these lease options:

- IP address of the client (required)
- Subnet mask of the client (required)
- DNS server IP address (optional)
- Router IP address (default gateway address to be used by the switch) (required)

If you want the switch to receive the configuration file from a TFTP server, you must configure the DHCP server with these lease options:

- TFTP server name (required)
- Boot filename (the name of the configuration file that the client needs) (recommended)
- Host name (optional)

Depending on the settings of the DHCP server, the switch can receive IP address information, the configuration file, or both.

If you do not configure the DHCP server with the lease options described earlier, it replies to client requests with only those parameters that are configured. If the IP address and subnet mask are not in the reply, the switch is not configured. If the router IP address or TFTP server name are not found, the switch might send broadcast, instead of unicast, TFTP requests. Unavailability of other lease options does not affect autoconfiguration.

The DHCP server can be on the same LAN or on a different LAN than the switch. If the DHCP server is running on a different LAN, you should configure a DHCP relay. For more information, see the “[Configuring the Relay Device](#)” section on [page 3-6](#). If your DHCP server is a Cisco device, refer to the “IP Addressing and Services” section in the *Cisco IOS IP and IP Routing Configuration Guide for Release 12.1*.

Configuring the TFTP Server

Based on the DHCP server configuration, the switch attempts to download one or more configuration files from the TFTP server. If you configured the DHCP server to respond to the switch with all the options required for IP connectivity to the TFTP server, and if you configured the DHCP server with a TFTP server name, address, and configuration filename, the switch attempts to download the specified configuration file from the specified TFTP server.

If you did not specify the configuration filename, the TFTP server, or if the configuration file could not be downloaded, the switch attempts to download a configuration file by using various combinations of filenames and TFTP server addresses. The files include the specified configuration filename (if any) and these files: `network-config`, `cisconet.cfg`, `hostname.config`, or `hostname.cfg`, where *hostname* is the switch's current hostname. The TFTP server addresses used include the specified TFTP server address (if any) and the broadcast address (255.255.255.255).

For the switch to successfully download a configuration file, the TFTP server must contain one or more configuration files in its base directory. The files can include these files:

- The configuration file named in the DHCP reply (the actual switch configuration file).
- The `network-config` or the `cisconet.cfg` file (known as the default configuration files).
- The `router-config` or the `ciscotr.cfg` file (These files contain commands common to all switches. Normally, if the DHCP and TFTP servers are properly configured, these files are not accessed.)

If you specify the TFTP server name in the DHCP server-lease database, you must also configure the TFTP server name-to-IP-address mapping in the DNS-server database.

If the TFTP server to be used is on a different LAN from the switch, or if it is to be accessed by the switch through the broadcast address (which occurs if the DHCP server response does not contain all the required information described earlier), a relay must be configured to forward the TFTP packets to the TFTP server. For more information, see the [“Configuring the Relay Device” section on page 3-6](#). The preferred solution is to configure the DHCP server with all the required information.

Configuring the DNS

The DHCP server uses the DNS server to resolve the TFTP server name to an IP address. You must configure the TFTP server name-to-IP address map on the DNS server. The TFTP server contains the configuration files for the switch.

You can configure the IP addresses of the DNS servers in the lease database of the DHCP server from where the DHCP replies will retrieve them. You can enter up to two DNS server IP addresses in the lease database.

The DNS server can be on the same or on a different LAN as the switch. If it is on a different LAN, the switch must be able to access it through a router.

Configuring the Relay Device

You must configure a relay device when a switch sends broadcast packets that need to be responded to by a host on a different LAN. Examples of broadcast packets that the switch might send are DHCP, DNS, and in some cases, TFTP packets. You must configure this relay device to forward received broadcast packets on an interface to the destination host.

If the relay device is a Cisco router or Layer 3 switch, enable IP routing (**ip routing** global configuration command), and configure helper addresses by using the **ip helper-address** interface configuration command.

For example, in [Figure 3-2](#), configure the router interfaces as follows:

On interface 10.0.0.2:

```
router(config-if)# ip helper-address 20.0.0.2
router(config-if)# ip helper-address 20.0.0.3
router(config-if)# ip helper-address 20.0.0.4
```

On interface 20.0.0.1


```
router(config-if)# ip helper-address 10.0.0.1
```



Note

If the switch is acting as the relay device, configure the interface as a routed port. For more information, see the [“Routed Ports” section on page 9-4](#) and the [“Configuring Layer 3 Interfaces” section on page 9-18](#).

Figure 3-2 Relay Device Used in Autoconfiguration



Obtaining Configuration Files

Depending on the availability of the IP address and the configuration filename in the DHCP reserved lease, the switch obtains its configuration information in these ways:

- The IP address and the configuration filename is reserved for the switch and provided in the DHCP reply (one-file read method).

The switch receives its IP address, subnet mask, TFTP server address, and the configuration filename from the DHCP server. The switch sends a unicast message to the TFTP server to retrieve the named configuration file from the base directory of the server, and upon receipt, completes its boot-up process.

- The IP address and the configuration filename is reserved for the switch, but the TFTP server address is not provided in the DHCP reply (one-file read method).

The switch receives its IP address, subnet mask, and the configuration filename from the DHCP server. The switch sends a broadcast message to a TFTP server to retrieve the named configuration file from the base directory of the server, and upon receipt, completes its boot-up process.

- Only the IP address is reserved for the switch and provided in the DHCP reply. The configuration filename is not provided (two-file read method).

The switch receives its IP address, subnet mask, and the TFTP server address from the DHCP server. The switch sends a unicast message to the TFTP server to retrieve the network-config or cisco.net.cfg default configuration file. (If the network-config file cannot be read, the switch reads the cisco.net.cfg file.)

The default configuration file contains the host names-to-IP-address mapping for the switch. The switch fills its host table with the information in the file and obtains its host name. If the host name is not found in the file, the switch uses the host name in the DHCP reply. If the host name is not specified in the DHCP reply, the switch uses the default *Switch* as its host name.

After obtaining its host name from the default configuration file or the DHCP reply, the switch reads the configuration file that has the same name as its host name (*hostname-conf* or *hostname.cfg*, depending on whether *network-conf* or *cisconet.cfg* was read earlier) from the TFTP server. If the *cisconet.cfg* file is read, the filename of the host is truncated to eight characters.

If the switch cannot read the *network-conf*, *cisconet.cfg*, or the *hostname* file, it reads the *router-conf* file. If the switch cannot read the *router-conf* file, it reads the *ciscortr.cfg* file.

**Note**

The switch broadcasts TFTP server requests if the TFTP server is not obtained from the DHCP replies, if all attempts to read the configuration file through unicast transmissions fail, or if the TFTP server name cannot be resolved to an IP address.

Example Configuration

Figure 3-3 shows a sample network for retrieving IP information by using DHCP-based autoconfiguration.

Figure 3-3 DHCP-Based Autoconfiguration Network Example




Table 3-2 shows the configuration of the reserved leases on the DHCP server.

Table 3-2 DHCP Server Configuration

	Switch A	Switch B	Switch C	Switch D
Binding key (hardware address)	00e0.9f1e.2001	00e0.9f1e.2002	00e0.9f1e.2003	00e0.9f1e.2004
IP address	10.0.0.21	10.0.0.22	10.0.0.23	10.0.0.24
Subnet mask	255.255.255.0	255.255.255.0	255.255.255.0	255.255.255.0
Router address	10.0.0.10	10.0.0.10	10.0.0.10	10.0.0.10
DNS server address	10.0.0.2	10.0.0.2	10.0.0.2	10.0.0.2
TFTP server name	maritsu or 10.0.0.3	maritsu or 10.0.0.3	maritsu or 10.0.0.3	maritsu or 10.0.0.3
Boot filename (configuration file) (optional)	switcha-conf	switchb-conf	switchc-conf	switchd-conf
Host name (optional)	switcha	switchb	switchc	switchd

DNS Server Configuration

The DNS server maps the TFTP server name *maritsu* to IP address 10.0.0.3.

TFTP Server Configuration (on UNIX)

The TFTP server base directory is set to `/tftpserver/work/`. This directory contains the `network-config` file used in the two-file read method. This file contains the host name to be assigned to the switch based on its IP address. The base directory also contains a configuration file for each switch (*switcha-config*, *switchb-config*, and so forth) as shown in this display:

```
prompt> cd /tftpserver/work/
prompt> ls
network-config
switcha-config
switchb-config
switchc-config
switchd-config
prompt> cat network-config
ip host switcha 10.0.0.21
ip host switchb 10.0.0.22
ip host switchc 10.0.0.23
ip host switchd 10.0.0.24
```

DHCP Client Configuration

No configuration file is present on Switch A through Switch D.

Configuration Explanation

In [Figure 3-3](#), Switch A reads its configuration file as follows:

- It obtains its IP address 10.0.0.21 from the DHCP server.
- If no configuration filename is given in the DHCP server reply, Switch A reads the `network-config` file from the base directory of the TFTP server.
- It adds the contents of the `network-config` file to its host table.
- It reads its host table by indexing its IP address 10.0.0.21 to its host name (switcha).
- It reads the configuration file that corresponds to its host name; for example, it reads *switcha-config* from the TFTP server.

Switches B through D retrieve their configuration files and IP addresses in the same way.

Manually Assigning IP Information

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

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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; do not enter leading zeros.
Step 3	ip address <i>ip-address subnet-mask</i>	Enter the IP address and subnet mask.
Step 4	exit	Return to global configuration mode.

	Command	Purpose
Step 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.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the switch IP address, use the **no ip address** interface configuration command. 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.

For information on setting the switch system name, protecting access to privileged EXEC commands, and setting time and calendar services, see [Chapter 5, “Administering the Switch.”](#)

Checking and Saving the Running Configuration

You can check the configuration settings you entered or changes you made by entering the **show running-config** privileged EXEC command: For information about the output of this command, refer to the *Cisco IOS Configuration Fundamental Command Reference for Release 12.1*.

To store the configuration or changes you have made to your startup configuration in flash memory, enter the **copy running-config startup-config** privileged EXEC command. 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.

For more information about alternative locations to copy the configuration file, see [Appendix B, “Working with the Cisco IOS File System, Configuration Files, and Software Images.”](#)

Modifying the Startup Configuration

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

- [Default Boot Configuration, page 3-11](#)
- [Automatically Downloading a Configuration File, page 3-11](#)
- [Booting Manually, page 3-12](#)
- [Booting a Specific Software Image, page 3-12](#)
- [Controlling Environment Variables, page 3-13](#)

Default Boot Configuration

Table 3-3 shows the default boot configuration.

Table 3-3 Default Boot Configuration

Feature	Default Setting
Operating system software image	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. The IOS image is stored in a directory that has the same name as the image file (excluding the .bin extension). In a depth-first search of a directory, each encountered subdirectory is completely searched before continuing the search in the original directory.
Configuration file	Configured switches use the <i>config.text</i> file stored on the system board in flash memory. A new switch has no configuration file.

Automatically Downloading a Configuration File

You can automatically download a configuration file to your switch by using the DHCP-based autoconfiguration feature. For more information, see the [“Understanding DHCP-Based Autoconfiguration”](#) section on page 3-3.

Specifying the Filename to Read and Write the System Configuration

By default, the IOS 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 is loaded during the next boot cycle.

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

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

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

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	boot manual	Enable the switch to manually boot during the next boot cycle.
Step 3	end	Return to privileged EXEC mode.
Step 4	show boot	Verify your entries. The boot manual global command changes the setting of the <code>MANUAL_BOOT</code> environment variable. The next time you reboot the system, the switch is in boot loader mode, shown by the <i>switch:</i> prompt. To boot the system, use the boot filesystem:/file-url boot loader command. - 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.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable manual booting, use the **no boot manual** global configuration command.

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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	boot system filesystem:/file-url	Configure the switch to boot a specific image in flash memory during the next boot cycle. - 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.

	Command	Purpose
Step 3	end	Return to privileged EXEC mode.
Step 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.
Step 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 switch console connection configured for 9600 bps. Unplug the switch power cord and press the switch **Mode** button while reconnecting the power cord. You can release the **Mode** button a second or two after the LED above port 1 turns off. Then the boot loader *switch:* prompt is displayed.

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 flash memory outside of the flash file system.

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 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 IOS configuration file. For example, the name of the 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 IOS 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 3-4 describes the function of the most common environment variables.

Table 3-4 Environment Variables

Variable	Boot Loader Command	IOS Global Configuration Command
BOOT	set BOOT <i>filesystem:/file-url ...</i> 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.	boot system <i>filesystem:/file-url</i> Specifies the IOS image to load during the next boot cycle. This command changes the setting of the BOOT environment variable.
MANUAL_BOOT	set MANUAL_BOOT <i>yes</i> Determines whether the switch automatically or manually boots. Valid values are 1, yes, 0, and no. If it is set to no or 0, the boot loader attempts to automatically boot the system. If it is set to anything else, you must manually boot the switch from the boot loader mode.	boot manual Enables manually booting the switch during the next boot cycle and changes the setting of the MANUAL_BOOT environment variable. The next time you reboot the system, the switch is in boot loader mode. To boot the system, use the boot loader boot flash: <i>filesystem:/file-url</i> command, and specify the name of the bootable image.
CONFIG_BUFSIZE	set CONFIG_BUFSIZE <i>size</i> Changes the buffer size that IOS 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 <i>size</i> 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.
CONFIG_FILE	set CONFIG_FILE flash: <i>/file-url</i> Changes the filename that IOS uses to read and write a nonvolatile copy of the system configuration.	boot config-file flash: <i>/file-url</i> Specifies the filename that IOS uses to read and write a nonvolatile copy of the system configuration. This command changes the CONFIG_FILE environment variable.

Scheduling a Reload of the Software Image

You can schedule a reload of the software image to occur on 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 2003 (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 2003 (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).



Configuring IE2100 CNS Agents

This chapter describes how to configure the Intelligence Engine 2100 (IE2100) Series Cisco Networking Services (CNS) embedded agents on your Catalyst 3750 Metro switch.



Note

For complete syntax and usage information for the commands used in this section, refer to the *Cisco Intelligence Engine 2100 Series Configuration Registrar Manual*, and select **Cisco IOS Software Release 12.2 > New Feature Documentation > 12.2(2)T** on Cisco.com.

This chapter consists of these sections:

- [Understanding IE2100 Series Configuration Registrar Software, page 4-1](#)
- [Understanding CNS Embedded Agents, page 4-5](#)
- [Configuring CNS Embedded Agents, page 4-6](#)
- [Displaying CNS Configuration, page 4-12](#)


Understanding IE2100 Series Configuration Registrar Software

The IE2100 Series Configuration Registrar is a network management device that acts as a configuration service for automating the deployment and management of network devices and services (see [Figure 4-1](#)). Each Configuration Registrar manages a group of Cisco IOS devices (switches and routers) and the services that they deliver, storing their configurations and delivering them as needed. The Configuration Registrar automates initial configurations and configuration updates by generating device-specific configuration changes, sending them to the device, executing the configuration change, and logging the results.

The Configuration Registrar supports standalone and server modes and has these CNS components:

- Configuration service (web server, file manager, and namespace mapping server)
- Event service (event gateway)
- Data service directory (data models and schema)

In standalone mode, the Configuration Registrar supports an embedded CNS Directory Service. In this mode, no external directory or other data store is required. In server mode, the Configuration Registrar supports the use of a user-defined external directory.

Figure 4-1 Configuration Registrar Architectural Overview

These sections contain this conceptual information:

- [CNS Configuration Service, page 4-2](#)
- [CNS Event Service, page 4-3](#)
- [Configuration Registrar and ConfigID, DeviceID, and Host Name, page 4-3](#)

CNS Configuration Service

The CNS Configuration Service is the core component of the Configuration Registrar. It consists of a configuration server that works with CNS configuration agents located on the switch. The CNS Configuration Service delivers device and service configurations to the switch for initial configuration and mass reconfiguration by logical groups. Switches receive their initial configuration from the CNS Configuration Service when they start up on the network for the first time.

The CNS Configuration Service uses the CNS Event Service to send and receive configuration change events and to send success and failure notifications.

The configuration server is a web server that uses configuration templates and the device-specific configuration information stored in the embedded (standalone mode) or remote (server mode) directory.

Configuration templates are text files containing static configuration information in the form of CLI commands. In the templates, variables are specified using lightweight directory access protocol (LDAP) URLs that reference the device-specific configuration information stored in a directory.

The configuration agent can perform a syntax check on received configuration files and publish events to indicate the success or failure of the syntax check. The configuration agent can either apply configurations immediately or delay the application until receipt of a synchronization event from the configuration server.

CNS Event Service

The Configuration Registrar uses the CNS Event Service for receipt and generation of configuration events. The CNS event agent resides on the switch and facilitates the communication between the switch and the event gateway on the Configuration Registrar.

The CNS Event Service is a highly scalable publish-and-subscribe communication method. The CNS Event Service uses subject-based addressing to send messages to their destinations. Subject-based addressing conventions define a simple, uniform namespace for messages and their destinations.

NameSpace Mapper

The Configuration Registrar includes the NameSpace Mapper (NSM) that provides a lookup service for managing logical groups of devices based on application, device ID or group ID, and event.

Cisco IOS devices recognize only event subject-names that match those configured in Cisco IOS software; for example, `cisco.cns.config.load`. You can use the namespace mapping service to designate events by using any desired naming convention. When you have populated your data store with your subject names, NSM resolves your event subject-name strings to those known by IOS.

For a subscriber, when given a unique device ID and event, the namespace mapping service returns a set of events to which to subscribe. Similarly, for a publisher, when given a unique group ID, device ID, and event, the mapping service returns a set of events on which to publish.

Configuration Registrar and ConfigID, DeviceID, and Host Name

The Configuration Registrar assumes that a unique identifier is associated with each configured switch. This unique identifier can take on multiple synonyms, where each synonym is unique within a particular namespace. The event service uses namespace content for subject-based addressing of messages.

The Configuration Registrar intersects two namespaces, one for the event bus and the other for the configuration server. Within the scope of the configuration server namespace, the term `configID` is the unique identifier for a device. Within the scope of the event bus namespace, the term `deviceID` is the CNS unique identifier for a device.

Because the Configuration Registrar uses both the event bus and the configuration server to provide configurations to devices, you must define both `configID` and `deviceID` for each configured switch.

Within the scope of a single instance of the configuration server, no two configured switches can share the same value for `configID`. Within the scope of a single instance of the event bus, no two configured switches can share the same value for `deviceID`.

ConfigID

Each configured switch has a unique `configID`, which serves as the key into the Configuration Registrar directory for the corresponding set of switch CLI attributes. The `configID` defined on the switch must match the `configID` for the corresponding switch definition on the Configuration Registrar.

The `configID` is fixed at boot time and cannot be changed until reboot, even when the switch host name is reconfigured.

DeviceID

Each configured switch participating on the event bus has a unique deviceID, which is analogous to the switch source address so that the switch can be targeted as a specific destination on the bus. All switches configured with the **cns config partial** global configuration command must access the event bus. Therefore, the deviceID, as originated on the switch, must match the deviceID of the corresponding switch definition in the Configuration Registrar.

The origin of the deviceID is defined by the Cisco IOS host name of the switch. However, the deviceID variable and its usage reside within the event gateway, which is adjacent to the switch.

The logical Cisco IOS termination point on the event bus is embedded in the event gateway, which in turn functions as a proxy on behalf of the switch. The event gateway represents the switch and its corresponding deviceID to the event bus.

The switch declares its host name to the event gateway immediately after the successful connection to the event gateway. The event gateway couples the deviceID value to the Cisco IOS host name each time this connection is established. The event gateway caches this deviceID value for the duration of its connection to the switch.

Host Name and DeviceID

The deviceID is fixed at the time of the connection to the event gateway and does not change even when the switch host name is reconfigured.

When changing the switch host name on the switch, the only way to refresh the deviceID is to break the connection between the switch and the event gateway. Enter the **no cns event** global configuration command followed by the **cns event** global configuration command.

When the connection is re-established, the switch sends its modified host name to the event gateway. The event gateway redefines the deviceID to the new value.



Caution

When using the Configuration Registrar user interface, you must first set the deviceID field to the host name value that the switch acquires *after*—not *before*—you use the **cns config initial** global configuration command at the switch. Otherwise, subsequent **cns config partial** global configuration command operations malfunction.

Using Host Name, DeviceID, and ConfigID

In standalone mode, when a host name value is set for a switch, the configuration server uses the host name as the deviceID when an event is sent on host name. If the host name has not been set, the event is sent on the cn=<value> of the device.

In server mode, the host name is not used. In this mode, the unique deviceID attribute is always used for sending an event on the bus. If this attribute is not set, you cannot update the switch.

These and other associated attributes (tag value pairs) are set when you run **Setup** on the Configuration Registrar.



Note

For more information about running the setup program on the Configuration Registrar, refer to the *Cisco Intelligence Engine 2100 Series Configuration Registrar Manual*.

Understanding CNS Embedded Agents

The CNS event agent feature allows the switch to publish and subscribe to events on the event bus and works with the CNS configuration agent. The CNS configuration agent feature supports the switch by providing:

- Initial configurations
- Incremental (partial) configurations
- Synchronized configuration updates

Initial Configuration


When the switch first comes up, it attempts to get an IP address by broadcasting a DHCP request on the network. Assuming there is no DHCP server on the subnet, the distribution switch acts as a DHCP relay agent and forwards the request to the DHCP server. Upon receiving the request, the DHCP server assigns an IP address to the new switch and includes the TFTP server IP address, the path to the bootstrap configuration file, and the default gateway IP address in a unicast reply to the DHCP relay agent. The DHCP relay agent forwards the reply to the switch.

The switch automatically configures the assigned IP address on interface VLAN 1 (the default) and downloads the bootstrap configuration file from the TFTP server. Upon successful download of the bootstrap configuration file, the switch loads the file in its running configuration.

The embedded CNS agents initiate communication with the IE2100 Configuration Registrar by using the appropriate configID and eventID. The Configuration Registrar maps the configID to a template and downloads the full configuration file to the switch.

Figure 4-2 shows a sample network configuration for retrieving the initial bootstrap configuration file by using DHCP-based autoconfiguration.

Figure 4-2 Initial Configuration Overview



Incremental (Partial) Configuration

After the network is running, you can add new services by using the CNS configuration agent. You can send incremental (partial) configurations to the switch. You can send the actual configuration as an event payload by way of the event gateway (push operation), or as a signal event that triggers the switch to initiate a pull operation.

The switch can check the syntax of the configuration before applying it. If the syntax is correct, the switch applies the incremental configuration and publishes an event that signals success to the configuration server. If the switch does not apply the incremental configuration, it publishes an event showing an error status. When the switch has applied the incremental configuration, it can write it to NVRAM or wait until signaled to do so.

Synchronized Configuration

When the switch receives a configuration, it can defer application of the configuration upon receipt of a write-signal event. The write-signal event tells the switch not to save the updated configuration into its NVRAM. The switch uses the updated configuration as its running configuration. This ensures that the switch configuration is synchronized with other network activities before saving the configuration in NVRAM for use at the next reboot.

Configuring CNS Embedded Agents

The CNS agents embedded in the switch Cisco IOS software allow the switch to be connected and automatically configured as described in the [“Enabling Automated CNS Configuration” section on page 4-6](#). If you want to change the configuration or install a custom configuration, see these sections for instructions:

- [Enabling the CNS Event Agent, page 4-8](#)
- [Enabling the CNS Configuration Agent, page 4-9](#)

Enabling Automated CNS Configuration

To enable automated CNS configuration of the switch, you must first complete the prerequisites in [Table 4-1](#). When you complete them, power on the switch. At the **setup** prompt, do nothing: The switch begins the initial configuration as described in the [“Initial Configuration” section on page 4-5](#). When the full configuration file is loaded on your switch, you need to do nothing else.

Table 4-1 Prerequisites for Enabling Automatic Configuration

Device	Required Configuration
Access switch	Factory default (no configuration file)
Distribution switch	• IP helper address • Enable DHCP relay agent • IP routing (if used as default gateway)

Table 4-1 Prerequisites for Enabling Automatic Configuration (continued)

Device	Required Configuration
DHCP server	• IP address assignment • TFTP server IP address • Path to bootstrap configuration file on the TFTP server • Default gateway IP address
TFTP server	• Create a bootstrap configuration file that includes the CNS configuration commands that enable the switch to communicate with the IE2100 Configuration Registrar. • Configure the switch to use either the switch MAC address or the serial number (instead of the default host name) to generate the configID and eventID. • Configure the CNS event agent to push the configuration file to the switch.
IE2100 Configuration Registrar	Create one or more templates for each type of device, and map the configID of the device to the template.

**Note**

For more information about running the setup program and creating templates on the Configuration Registrar, refer to the *Cisco Intelligence Engine 2100 Series Configuration Registrar Manual*.

Enabling the CNS Event Agent


Note

You must enable the CNS event agent on the switch before you enable the CNS configuration agent.

Beginning in privileged EXEC mode, follow these steps to enable the CNS event agent on the switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	cns event { <i>ip-address</i> <i>hostname</i> } [<i>port-number</i>] [backup] [init-retry <i>retry-count</i>] [keepalive <i>seconds</i> <i>retry-count</i>] [source <i>ip-address</i>]	Enable the event agent, and enter the gateway parameters. - For {<i>ip-address</i> <i>hostname</i>}, enter either the IP address or the host name of the event gateway. (Optional) For <i>port number</i>, enter the port number for the event gateway. The default port number is 11011. (Optional) Enter backup to show that this is the backup gateway. (If omitted, this is the primary gateway.) (Optional) For init-retry <i>retry-count</i>, enter the number of initial retries before switching to backup. The default is 3. (Optional) For keepalive <i>seconds</i>, enter how often the switch sends keepalive messages. For <i>retry-count</i>, enter the number of unanswered keepalive messages that the switch sends before the connection is terminated. The default for each is 0. (Optional) For source <i>ip-address</i>, enter the source IP address of this device. Note Though visible in the command-line help string, the encrypt and force-fmt1 keywords are not supported.
Step 3	end	Return to privileged EXEC mode.
Step 4	show cns event connections	Verify information about the event agent.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the CNS event agent, use the **no cns event** {*ip-address* | *hostname*} global configuration command.

This example shows how to enable the CNS event agent, set the IP address gateway to 10.180.1.27, set 120 seconds as the keepalive interval, and set 10 as the retry count.

```
Switch(config)# cns event 10.180.1.27 keepalive 120 10
```

Enabling the CNS Configuration Agent

After enabling the CNS event agent, start the CNS configuration agent on the switch. You can enable the configuration agent with these commands:

- The **cns config initial** global configuration command enables the configuration agent and initiates an initial configuration on the switch.
- The **cns config partial** global configuration command enables the configuration agent and initiates a partial configuration on the switch. You can then remotely send incremental configurations to the switch from the Configuration Registrar.

Enabling an Initial Configuration

Beginning in privileged EXEC mode, follow these steps to enable the CNS configuration agent and initiate an initial configuration on the switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	cns config connect-intf <i>interface-prefix</i> [ping-interval <i>seconds</i>] [retries <i>num</i>]	Enter the connect-interface-config mode, and specify the interface for connecting to the Configuration Registrar. • Enter the <i>interface-prefix</i> for the connecting interface. You must specify the interface type but need not specify the interface number. • (Optional) For ping-interval <i>seconds</i>, enter the interval between successive ping attempts. The range is 1 to 30 seconds. The default is 10 seconds. • (Optional) For retries <i>num</i>, enter the number of ping retries. The range is 1 to 30. The default is 5.
Step 3	config-cli or line-cli	Enter config-cli to connect to the Configuration Registrar through the interface defined in cns config connect-intf . Enter line-cli to connect to the Registrar through modem dialup lines. Note The config-cli interface configuration command accepts the special directive character & that acts as a placeholder for the interface name. When the configuration is applied, the & is replaced with the interface name. For example, to connect through FastEthernet1/0/1, the command config-cli ip route 0.0.0.0 0.0.0.0 & generates the command ip route 0.0.0.0 0.0.0.0 FastEthernet1/0/1.
Step 4	exit	Return to global configuration mode.
Step 5	hostname <i>name</i>	Enter the host name for the switch.
Step 6	ip route <i>network-number</i>	Establish a static route to the Configuration Registrar whose IP address is <i>network-number</i> .

	Command	Purpose
Step 7	cns id <i>interface num</i> { dns-reverse ipaddress mac-address } [event] or cns id { hardware-serial hostname string string } [event]	Set the unique eventID or configID used by the Configuration Registrar. - For <i>interface num</i>, enter the type of interface—for example, Ethernet, Group-Async, Loopback, or Virtual-Template. This setting specifies from which interface the IP or MAC address should be retrieved to define the unique ID. - For {dns-reverse ipaddress mac-address} enter dns-reverse to retrieve the host name and assign it as the unique ID, enter ipaddress to use the IP address, or enter mac-address to use the MAC address as the unique ID. (Optional) Enter event to set the ID to be the event-id value used to identify the switch. - For {hardware-serial hostname string string}, enter hardware-serial to set the switch serial number as the unique ID, enter hostname (the default) to select the switch host name as the unique ID, or enter an arbitrary text string for string string as the unique ID.
Step 8	cns config initial { <i>ip-address</i> <i>hostname</i> } [<i>port-number</i>] [event] [no-persist] [page page] [source ip-address] [syntax-check]	Enable the configuration agent, and initiate an initial configuration. - For {<i>ip-address</i> <i>hostname</i>}, enter the IP address or the host name of the configuration server. (Optional) For <i>port-number</i>, enter the port number of the configuration server. The default port number is 80. (Optional) Enable event for configuration success, failure, or warning messages when the configuration is finished. (Optional) Enable no-persist to suppress the automatic writing to NVRAM of the configuration pulled as a result of entering the cns config initial global configuration command. If the no-persist keyword is not entered, using the cns config initial command causes the resultant configuration to be automatically written to NVRAM. (Optional) For page page, enter the web page of the initial configuration. The default is /Config/config/asp. (Optional) Enter source ip-address to use for source IP address. (Optional) Enable syntax-check to check the syntax when this parameter is entered. Note Though visible in the command-line help string, the encrypt keyword is not supported.
Step 9	end	Return to privileged EXEC mode.
Step 10	show cns config connections	Verify information about the configuration agent.
Step 11	show running-config	Verify your entries.

To disable the CNS configuration agent, use the **no cns config initial** {*ip-address* | *hostname*} global configuration command.

This example shows how to configure an initial configuration on a remote switch. The switch host name is the unique ID. The CNS Configuration Registrar IP address is 172.28.129.22.

```
Switch(config)# cns config connect-intf serial ping-interval 1 retries 1
Switch(config-cns-conn-if)# config-cli ip address negotiated
Switch(config-cns-conn-if)# config-cli encapsulation ppp
Switch(config-cns-conn-if)# config-cli ip directed-broadcast
Switch(config-cns-conn-if)# config-cli no keepalive
Switch(config-cns-conn-if)# config-cli no shutdown
Switch(config-cns-conn-if)# exit
Switch(config)# hostname RemoteSwitch
RemoteSwitch(config)# ip route 10.1.1.1 255.255.255.255 11.11.11.1
RemoteSwitch(config)# cns id Ethernet 0 ipaddress
RemoteSwitch(config)# cns config initial 10.1.1.1 no-persist
```

Enabling a Partial Configuration

Beginning in privileged EXEC mode, follow these steps to enable the CNS configuration agent and to initiate a partial configuration on the switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	cns config partial { <i>ip-address</i> <i>hostname</i> } [<i>port-number</i>] [source <i>ip-address</i>]	Enable the configuration agent, and initiate a partial configuration. - For {<i>ip-address</i> <i>hostname</i>}, enter the IP address or the host name of the configuration server. (Optional) For <i>port-number</i>, enter the port number of the configuration server. The default port number is 80. (Optional) Enter source <i>ip-address</i> to use for the source IP address. Note Though visible in the command-line help string, the encrypt keyword is not supported.
Step 3	end	Return to privileged EXEC mode.
Step 4	show cns config stats or show cns config outstanding	Verify information about the configuration agent.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the CNS configuration agent, use the **no cns config partial** {*ip-address* | *hostname*} global configuration command. To cancel a partial configuration, use the **cns config cancel** privileged EXEC command.

Displaying CNS Configuration

You can use the privileged EXEC commands in [Table 4-2](#) to display CNS configuration information.

Table 4-2 *Displaying CNS Configuration*

Command	Purpose
show cns config connections	Display the status of the CNS configuration agent connections.
show cns config outstanding	Display information about incremental (partial) CNS configurations that have started but are not yet completed.
show cns config stats	Display statistics about the CNS configuration agent.
show cns event connections	Display the status of the CNS event agent connections.
show cns event stats	Display statistics about the CNS event agent.
show cns event subject	Display a list of event agent subjects that are subscribed to by applications.



Administering the Switch

This chapter describes how to perform one-time operations to administer the Catalyst 3750 Metro switch.

This chapter consists of these sections:

- [Managing the System Time and Date, page 5-1](#)
- [Configuring a System Name and Prompt, page 5-15](#)
- [Creating a Banner, page 5-18](#)
- [Managing the MAC Address Table, page 5-20](#)
- [Managing the ARP Table, page 5-26](#)

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 IOS Configuration Fundamentals Command Reference for Release 12.1*.

This section contains this configuration information:

- [Understanding the System Clock, page 5-2](#)
- [Understanding Network Time Protocol, page 5-2](#)
- [Configuring NTP, page 5-4](#)
- [Configuring Time and Date Manually, page 5-11](#)

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 5-11](#).

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 5-1 shows a typical network example using NTP. Switch A is the NTP master, with Switches B, C, and D configured in NTP server mode, in server association with Switch A. Switch E is configured as an NTP peer to the upstream and downstream switches, Switch B and Switch F.

Figure 5-1 Typical NTP Network Configuration



101349

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.

Configuring NTP

The switches do not have a hardware-supported clock, and they cannot function as an NTP master clock to which peers synchronize themselves when an external NTP source is not available. These switches also have 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, page 5-4](#)
- [Configuring NTP Authentication, page 5-5](#)
- [Configuring NTP Associations, page 5-6](#)
- [Configuring NTP Broadcast Service, page 5-7](#)
- [Configuring NTP Access Restrictions, page 5-8](#)
- [Configuring the Source IP Address for NTP Packets, page 5-10](#)
- [Displaying the NTP Configuration, page 5-11](#)

Default NTP Configuration

[Table 5-1](#) shows the default NTP configuration.

Table 5-1 *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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ntp authenticate	Enable the NTP authentication feature, which is disabled by default.
Step 3	ntp authentication-key <i>number</i> md5 <i>value</i>	Define the authentication keys. By default, none are defined. - 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. 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.
Step 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. By default, no trusted keys are defined. For <i>key-number</i>, specify the key defined in Step 3. This command provides protection against accidentally synchronizing the switch to a device that is not trusted.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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. - 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 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 provides 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to send NTP broadcast packets, and enter interface configuration mode.
Step 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. (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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.
Step 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 a port to send NTP version 2 packets:

```
Switch(config)# interface gigabitethernet1/0/1
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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to receive NTP broadcast packets, and enter interface configuration mode.
Step 3	ntp broadcast client	Enable the interface to receive NTP broadcast packets. By default, no interfaces receive NTP broadcast packets.
Step 4	exit	Return to global configuration mode.
Step 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.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable a port 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 a port to receive NTP broadcast packets:

```
Switch(config)# interface gigabitethernet1/0/1
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, page 5-9](#)
- [Disabling NTP Services on a Specific Interface, page 5-10](#)

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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ntp access-group {query-only serve-only serve peer} access-list-number	Create an access group, and apply a basic IP access list. The keywords have these meanings: • 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. For <i>access-list-number</i>, enter a standard IP access list number from 1 to 99.
Step 3	access-list access-list-number permit source [source-wildcard]	Create the access list. • 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. 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 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:

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

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 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 5-6.

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 IOS Configuration Fundamentals Command Reference for Release 12.1*.

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, page 5-11](#)
- [Displaying the Time and Date Configuration, page 5-12](#)
- [Configuring the Time Zone, page 5-12](#)
- [Configuring Summer Time \(Daylight Saving Time\), page 5-13](#)

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:

	Command	Purpose
Step 1	clock set hh:mm:ss day month year or clock set hh:mm:ss month day year	Manually set the system clock using one of these formats. • 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).
Step 2	show running-config	Verify your entries.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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. • 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	clock summer-time zone recurring [<i>week day month hh:mm week day month hh:mm [offset]</i>]	Configure summer time to start and end on the specified days every year. 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. - 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 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):

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	clock summer-time zone date [<i>month date year hh:mm month date year hh:mm [offset]</i>] or clock summer-time zone date [<i>date month year hh:mm date month year hh:mm [offset]</i>]	Configure summer time to start on the first date and end on the second date. Summer time is disabled by default. - 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 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, 2003, at 02:00, and end on April 26, 2004, at 02:00:

```
Switch(config)# clock summer-time pdt date 12 October 2003 2:00 26 April 2004 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 IOS Configuration Fundamentals Command Reference* and the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

This section contains this configuration information:

- [Default System Name and Prompt Configuration, page 5-15](#)
- [Configuring a System Name, page 5-15](#)
- [Configuring a System Prompt, page 5-16](#)
- [Understanding DNS, page 5-16](#)

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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 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 *cisco.com*. A specific device in this domain, for example, the FTP system is identified as *ftp.cisco.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, page 5-17](#)
- [Setting Up DNS, page 5-17](#)
- [Displaying the DNS Configuration, page 5-18](#)

Default DNS Configuration

Table 5-2 shows the default DNS configuration.

Table 5-2 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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; however, if the switch configuration comes from a BOOTP or DHCP server, then the default domain name might be set by the BOOTP or DHCP server (if the servers were configured with this information).
Step 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.
Step 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).
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 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 Cisco IOS 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 IOS Configuration Fundamentals Command Reference for Release 12.1*.

This section contains this configuration information:

- [Default Banner Configuration, page 5-18](#)
- [Configuring a Message-of-the-Day Login Banner, page 5-19](#)
- [Configuring a Login Banner, page 5-20](#)

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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	banner motd <i>c message c</i>	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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	banner login <i>c message c</i>	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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 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 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 and the type (static or dynamic).



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, page 5-21](#)
- [MAC Addresses and VLANs, page 5-21](#)
- [Default MAC Address Table Configuration, page 5-21](#)
- [Changing the Address Aging Time, page 5-22](#)
- [Removing Dynamic Address Entries, page 5-22](#)
- [Configuring MAC Address Notification Traps, page 5-23](#)
- [Adding and Removing Static Address Entries, page 5-24](#)
- [Displaying Address Table Entries, page 5-25](#)

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 globally configured. 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 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. Unicast addresses, for example, could forward to port 1 in VLAN 1 and ports 9, 10, and 1 in VLAN 5.



Note

Multiport static addresses are not supported.

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.

Default MAC Address Table Configuration

[Table 5-3](#) shows the default MAC address table configuration.

Table 5-3 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. You can change the aging time setting for all VLANs or for a specified VLAN.

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. Flooding results, which can impact switch performance.

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

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mac address-table aging-time [0 10-1000000] [vlan <i>vlan-id</i>]	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. For <i>vlan-id</i> , the range is 1 to 4094.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mac address-table aging-time	Verify your entries.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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. - 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.
Step 3	snmp-server enable traps mac-notification	Enable the switch to send MAC address traps to the NMS.
Step 4	mac address-table notification	Enable the MAC address notification feature.
Step 5	mac address-table notification [interval <i>value</i>] [history-size <i>value</i>]	Enter the trap interval time and the history table size. (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.

	Command	Purpose
Step 6	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 2 interface on which to enable the SNMP MAC address notification trap.
Step 7	snmp trap mac-notification { added removed }	Enable the MAC address notification trap. • 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.
Step 8	end	Return to privileged EXEC mode.
Step 9	show mac address-table notification interface show running-config	Verify your entries.
Step 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 a port.

```
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 fastethernet1/0/4
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 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.

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 unicast address 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mac address-table static <i>mac-addr</i> vlan <i>vlan-id</i> interface <i>interface-id</i>	Add a static address to the MAC address table. - For <i>mac-addr</i>, specify the destination MAC unicast address 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. The range is 1 to 4094. - For <i>interface-id</i>, specify the interface to which the received packet is forwarded. Valid interfaces include physical ports.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mac address-table static	Verify your entries.
Step 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 packet is forwarded to the specified interface:

```
Switch(config)# mac address-table static c2f3.220a.12f4 vlan 4 interface
gigabitethernet1/0/1
```

Displaying Address Table Entries

You can display the MAC address table by using one or more of the privileged EXEC commands described in [Table 5-4](#):

Table 5-4 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 only dynamic MAC address table entries.
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.
show mac address-table notification	Displays the MAC notification parameters and history table.

Table 5-4 Commands for Displaying the MAC Address Table (continued)

Command	Description
<code>show mac address-table static</code>	Displays only static MAC address table entries.
<code>show mac address-table vlan</code>	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.

For more information, refer to the Cisco IOS Release 12.1 documentation on Cisco.com.



Configuring SDM Templates

This chapter describes how to configure the Switch Database Management (SDM) templates on the Catalyst 3750 Metro 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 the SDM Templates, page 6-1](#)
- [Configuring the Switch SDM Template, page 6-2](#)
- [Displaying the SDM Templates, page 6-4](#)

Understanding the SDM Templates

You can use SDM templates to configure system resources in the switch to optimize support for specific features, depending on how the switch is used in the network. You can select a template to provide maximum system usage for some functions or to use the default template to balance resources.

The templates prioritize system resources to optimize support for these types of features:

- **Routing**—The routing template maximizes system resources for unicast routing, typically required for a router or aggregator in the center of a network. This template does not support policy-based routing.
- **Policy-based routing (PBR)**—The PBR template maximizes system resources to support PBR. This template is not compatible with virtual private network (VPN) routing/forwarding tables (VRFs).
- **VLANs**—The VLAN template disables routing and supports the maximum number of unicast MAC addresses. It would typically be selected for a Layer 2 switch.
- **Default**—The default template gives balance to all functions.

Table 6-1 lists the approximate numbers of each resource supported in each of the four templates.

**Note**

Although the resources allocated to features are the same in the routing and PBR templates, you must use the PBR template when enabling PBR, and you cannot use the PBR template when VRFs are configured on the switch.

Table 6-1 *Approximate Number of Feature Resources Allowed by Each Template*

Resource	Template			
	Default	Routing	PBR	VLAN
Unicast MAC addresses	6 K	3 K	3 K	12 K
IGMP groups and multicast routes	1 K	1 K	1 K	1 K
Unicast routes	8 K	11 K	11 K	0
• Directly connected hosts	6 K	3 K	3 K	0
• Indirect routes	2 K	8 K	8 K	0
Policy-based routing ACEs	0	512	512	0
QoS classification ACEs	512	512	512	512
Security ACEs	1 K	1 K	1 K	1 K
Layer 2 VLANs	1 K	1 K	1 K	1 K

The first eight rows in the tables (unicast MAC addresses through security ACEs) represent approximate hardware boundaries set when a template is selected. If a section of a hardware resource is full, all processing overflow is sent to the CPU, seriously impacting switch performance. The last row is a guideline used to calculate hardware resource consumption related to the number of Layer 2 VLANs on the switch.

Configuring the Switch SDM Template

This section describes how to configure the SDM template to be used on the switch. This section contains this configuration information:

- [Default SDM Template, page 6-2](#)
- [SDM Template Configuration Guidelines, page 6-3](#)
- [Setting the SDM Template, page 6-3](#)

Default SDM Template

The default template for is the “default desktop” template.

SDM Template Configuration Guidelines

You must reload the switch for the configuration to take effect.

Use the **sdm prefer vlan** global configuration command only on switches intended for Layer 2 switching with no routing. When you use the VLAN template, no system resources are reserved for routing entries, and any routing is done through software. This overloads the CPU and severely degrades routing performance.

Do not use the routing template if you do not have routing enabled on your switch. The **sdm prefer routing** global configuration command prevents other features from using the memory allocated to unicast routing in the routing template.

You must enter the **sdm prefer routing-pbr** global configuration command for PBR to function correctly.

VRFs are not compatible with the PBR template. Therefore multiprotocol label switching virtual private networks (MPLS VPNs) and multiple VRFs in CE devices (multi-VRF CE) are not compatible with PBR. When you enter the **sdm prefer routing-pbr** command, any preconfigured VRFs are removed from the configuration.

Setting the SDM Template

Beginning in privileged EXEC mode, follow these steps to use the SDM template to maximize resources:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	sdm prefer {default routing routing-pbr vlan}	Specify the SDM template to be used on the switch: The keywords have these meanings: • default—Sets the switch to the balance system resources between Layer 2 and Layer 3 functions. • routing—Maximizes routing on the switch. • routing-pbr—Allows PBR on the switch. • vlan—Maximizes VLAN configuration on the switch with no routing supported in hardware.
Step 3	end	Return to privileged EXEC mode.
Step 4	reload	Reload the operating system.

After the system reboots, you can use the **show sdm prefer** privileged EXEC command to verify the new template configuration. If you enter the **show sdm prefer** command before you enter the **reload** privileged EXEC command, the **show sdm prefer** command shows the template currently in use and the template that will become active after a reload.

This is an example of an output display when you have changed the template and have not reloaded the switch:

```
Switch# show sdm prefer
The current template is "desktop routing" template.
The selected template optimizes the resources in
the switch to support this level of features for
8 routed interfaces and 1024 VLANs.
```

```

number of unicast mac addresses:      3K
number of igmp groups + multicast routes: 1K
number of unicast routes:            11K
  number of directly connected hosts:  3K
  number of indirect routes:           8K
number of qos aces:                   512
number of security aces:              1K

```

On next reload, template will be "desktop vlan" template.

To return to the default template, use the **no sdm prefer** global configuration command.

This example shows how to configure a switch with the routing template.

```

Switch(config)# sdm prefer routing
Switch(config)# end
Switch# reload
Proceed with reload? [confirm]

```

Displaying the SDM Templates

Use the **show sdm prefer** privileged EXEC command with no parameters to display the active template.

Use the **show sdm prefer [default | routing | routing-pbr | vlan]** privileged EXEC command to display the resource numbers supported by the specified template.

This is an example of output from the **show sdm prefer** command, displaying the template in use.

```

Switch# show sdm prefer
The current template is "desktop default" template.
The selected template optimizes the resources in
the switch to support this level of features for
8 routed interfaces and 1024 VLANs.

```

```

number of unicast mac addresses:      6K
number of igmp groups + multicast routes: 1K
number of unicast routes:            8K
  number of directly connected hosts:  6K
  number of indirect routes:           2K
number of policy based routing aces:   0
number of qos aces:                   512
number of security aces:              1K

```

This is an example of output from the **show sdm prefer routing** command:

```

Switch# show sdm prefer routing
"desktop routing" template:
The selected template optimizes the resources in
the switch to support this level of features for
8 routed interfaces and 1024 VLANs.

number of unicast mac addresses:      3K
number of igmp groups + multicast routes: 1K
number of unicast routes:            11K
  number of directly connected hosts:  3K
  number of indirect routes:           8K
number of policy based routing aces:   512
number of qos aces:                   512
number of security aces:              1K

```



Configuring Switch-Based Authentication

This chapter describes how to configure switch-based authentication on the Catalyst 3750 Metro switch. It consists of these sections:

- [Preventing Unauthorized Access to Your Switch, page 7-1](#)
- [Protecting Access to Privileged EXEC Commands, page 7-2](#)
- [Controlling Switch Access with TACACS+, page 7-10](#)
- [Controlling Switch Access with RADIUS, page 7-17](#)
- [Controlling Switch Access with Kerberos, page 7-31](#)
- [Configuring the Switch for Local Authentication and Authorization, page 7-36](#)
- [Configuring the Switch for Secure Shell, page 7-37](#)

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 7-2](#).
- 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 ports 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 7-6](#).
- 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 7-10](#).

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 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, page 7-2](#)
- [Setting or Changing a Static Enable Password, page 7-3](#)
- [Protecting Enable and Enable Secret Passwords with Encryption, page 7-3](#)
- [Disabling Password Recovery, page 7-5](#)
- [Setting a Telnet Password for a Terminal Line, page 7-6](#)
- [Configuring Username and Password Pairs, page 7-6](#)
- [Configuring Multiple Privilege Levels, page 7-8](#)

Default Password and Privilege Level Configuration

[Table 7-1](#) shows the default password and privilege level configuration.

Table 7-1 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	enable password <i>password</i>	Define a new password or change an existing password for access to privileged EXEC mode. By default, no password is defined. 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 Ctrl-v when you create the password; for example, to create the password abc?123, do this: Enter abc. Enter Ctrl-v. Enter ?123. 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 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 *l1u2c3k4y5*. The password is not encrypted and provides access to level 15 (traditional privileged EXEC mode access):

```
Switch(config)# enable password l1u2c3k4y5
```

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 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	enable password [level <i>level</i>] { <i>password</i> <i>encryption-type encrypted-password</i> } or enable secret [level <i>level</i>] { <i>password</i> <i>encryption-type 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. (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 that you copy from another switch configuration. 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.
Step 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.
Step 4	end	Return to privileged EXEC mode.
Step 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 7-8](#).

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

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 `$1$FaD0$Xyti5Rkls3LoyxzS8` for privilege level 2:

```
Switch(config)# enable secret level 2 5 $1$FaD0$Xyti5Rkls3LoyxzS8
```

Disabling Password Recovery

By default, any end user with physical access to the switch can recover from a lost password by interrupting the boot process while the switch is powering on and then by entering a new password.

The password-recovery disable feature protects access to the switch password by disabling part of this functionality. When this feature is enabled, the end user can interrupt the boot process only by agreeing to set the system back to the default configuration. With password recovery disabled, you can still interrupt the boot process and change the password, but the configuration file (config.text) and the VLAN database file (vlan.dat) are deleted.



Note

If you disable password recovery, we recommend that you keep a backup copy of the configuration file on a secure server in case the end user interrupts the boot process and sets the system back to default values. Do not keep a backup copy of the configuration file on the switch. If the switch is operating in VTP transparent mode, we recommend that you also keep a backup copy of the VLAN database file on a secure server. When the switch is returned to the default system configuration, you can download the saved files to the switch by using the Xmodem protocol. For more information, see the [“Recovering from a Lost or Forgotten Password”](#) section on page 34-3.

Beginning in privileged EXEC mode, follow these steps to disable password recovery:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no service password-recovery	Disable password recovery. This setting is saved in an area of the flash memory that is accessible by the boot loader and the Cisco IOS image, but it is not part of the file system and is not accessible by any user.
Step 3	end	Return to privileged EXEC mode.
Step 4	show version	Verify the configuration by checking the last few lines of the display.

To re-enable password recovery, use the **service password-recovery** global configuration command.



Note

Disabling password recovery does not work if you have set the switch to boot manually by using the **boot manual** global configuration command. This command produces the boot loader prompt (*switch:~*) after the switch is power cycled.

Setting a Telnet Password for a Terminal Line

When you power-up your switch for the first time, an automatic setup program runs to assign IP information and to create a default configuration for continued use. The setup program also prompts you to configure your switch for Telnet access through a password. If you did not configure this password during the setup program, you can configure it now through the command-line interface (CLI).

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

	Command	Purpose
Step 1		Attach a PC or workstation with emulation software to the switch console port. The default data characteristics of the console port are 9600, 8, 1, no parity. You might need to press the Return key several times to see the command-line prompt.
Step 2	enable password <i>password</i>	Enter privileged EXEC mode.
Step 3	configure terminal	Enter global configuration mode.
Step 4	line vty 0 15	Configure the number of Telnet sessions (lines), and enter line configuration mode. There are 16 possible sessions on a command-capable switch. The 0 and 15 mean that you are configuring all 16 possible Telnet sessions.
Step 5	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.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries. The password is listed under the command line vty 0 15 .
Step 8	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 ports 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	username <i>name</i> [privilege <i>level</i>] { password <i>encryption-type password</i> }	Enter the username, privilege level, and password for each user. - 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.
Step 3	line console 0 or line vty 0 15	Enter line configuration mode, and configure the console port (line 0) or the VTY lines (line 0 to 15).
Step 4	login local	Enable local password checking at login time. Authentication is based on the username specified in Step 2.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 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, page 7-8](#)
- [Changing the Default Privilege Level for Lines, page 7-9](#)
- [Logging into and Exiting a Privilege Level, page 7-10](#)

Setting the Privilege Level for a Command

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

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	privilege mode level level command	Set the privilege level for a command. • 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.
Step 3	enable password level level password	Specify the enable password for the privilege level. • 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.
Step 4	end	Return to privileged EXEC mode.
Step 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.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	line vty line	Select the virtual terminal line on which to restrict access.
Step 3	privilege level level	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.
Step 4	end	Return to privileged EXEC mode.
Step 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.
Step 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. You might specify a high level or privilege level for your console line to restrict line usage.

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:

	Command	Purpose
Step 1	enable <i>level</i>	Log in to a specified privilege level. For <i>level</i> , the range is 0 to 15.
Step 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 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 Release 12.1*.

This section contains this configuration information:


- [Understanding TACACS+, page 7-10](#)
- [TACACS+ Operation, page 7-12](#)
- [Configuring TACACS+, page 7-12](#)
- [Displaying the TACACS+ Configuration, page 7-17](#)

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 7-1](#).

Figure 7-1 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 and 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, page 7-13](#)
- [Identifying the TACACS+ Server Host and Setting the Authentication Key, page 7-13](#)
- [Configuring TACACS+ Login Authentication, page 7-14](#)

- [Configuring TACACS+ Authorization for Privileged EXEC Access and Network Services, page 7-16](#)
- [Starting TACACS+ Accounting, page 7-16](#)

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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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. • 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.
Step 3	aaa new-model	Enable AAA.
Step 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.
Step 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.

	Command	Purpose
Step 6	end	Return to privileged EXEC mode.
Step 7	show tacacs	Verify your entries.
Step 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 ports. 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 port 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 ports 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa new-model	Enable AAA.

	Command	Purpose
Step 3	aaa authentication login { default <i>list-name</i> } <i>method1</i> [<i>method2</i> ...]	Create a login authentication method list. - 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 ports. - 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. Select one of these methods: 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 7-13. 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.
Step 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 5	login authentication { default <i>list-name</i> }	Apply the authentication list to a line or set of lines. - 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.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa authorization network tacacs+	Configure the switch for user TACACS+ authorization for all network-related service requests.
Step 3	aaa authorization exec tacacs+	Configure the switch for user TACACS+ authorization if the user has privileged EXEC access. The exec keyword might return user profile information (such as autocommand information).
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa accounting network start-stop tacacs+	Enable TACACS+ accounting for all network-related service requests.
Step 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 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 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 Release 12.1*.

This section contains this configuration information:

- [Understanding RADIUS, page 7-18](#)
- [RADIUS Operation, page 7-19](#)
- [Configuring RADIUS, page 7-20](#)
- [Displaying the RADIUS Configuration, page 7-31](#)

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 7-2 on page 7-19](#).
- 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 7-2 Transitioning from RADIUS to TACACS+ Services

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:
 - a. **ACCEPT**—The user is authenticated.
 - b. **REJECT**—The user is either not authenticated and is prompted to re-enter the username and password, or access is denied.
 - c. **CHALLENGE**—A challenge requires additional data from the user.
 - d. **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, page 7-20](#)
- [Identifying the RADIUS Server Host, page 7-20](#) (required)
- [Configuring RADIUS Login Authentication, page 7-23](#) (required)
- [Defining AAA Server Groups, page 7-25](#) (optional)
- [Configuring RADIUS Authorization for User Privileged Access and Network Services, page 7-27](#) (optional)
- [Starting RADIUS Accounting, page 7-28](#) (optional)
- [Configuring Settings for All RADIUS Servers, page 7-28](#) (optional)
- [Configuring the Switch to Use Vendor-Specific RADIUS Attributes, page 7-29](#) (optional)
- [Configuring the Switch for Vendor-Proprietary RADIUS Server Communication, page 7-30](#) (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 7-28](#).

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 7-25](#).

Beginning in privileged EXEC mode, follow these steps to configure per-server RADIUS server communication. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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>]	Specify the IP address or host name of the remote RADIUS server host. • (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. 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. 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 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 ports. 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 port 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 ports 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa new-model	Enable AAA.

	Command	Purpose
Step 3	aaa authentication login { default <i>list-name</i> } <i>method1</i> [<i>method2...</i>]	Create a login authentication method list. - 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 ports. - 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. Select one of these methods: 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 7-20. 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.
Step 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.

	Command	Purpose
Step 5	login authentication { default <i>list-name</i> }	Apply the authentication list to a line or set of lines. - 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.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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>]	Specify the IP address or host name of the remote RADIUS server host. (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. 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. 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.
Step 3	aaa new-model	Enable AAA.
Step 4	aaa group server radius <i>group-name</i>	Define the AAA server-group with a group name. This command puts the switch in a server group configuration mode.
Step 5	server <i>ip-address</i>	Associate a particular RADIUS server with the defined server group. Repeat this step for each RADIUS server in the AAA server group. Each server in the group must be previously defined in Step 2.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.
Step 9		Enable RADIUS login authentication. See the “Configuring RADIUS Login Authentication” section on page 7-23.

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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa authorization network radius	Configure the switch for user RADIUS authorization for all network-related service requests.
Step 3	aaa authorization exec radius	Configure the switch for user RADIUS authorization if the user has privileged EXEC access. The exec keyword might return user profile information (such as autocommand information).
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa accounting network start-stop radius	Enable RADIUS accounting for all network-related service requests.
Step 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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.
Step 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 1 to 1000.
Step 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.

	Command	Purpose
Step 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.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your settings.
Step 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 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"
```

This example shows how to apply an input ACL in ASCII format to a port for the duration of this connection:

```
cisco-avpair= "ip:inacl#1=deny ip 10.10.10.10 0.0.255.255 20.20.20.20 255.255.0.0"
cisco-avpair= "ip:inacl#2=deny ip 10.10.10.10 0.0.255.255 any"
cisco-avpair= "mac:inacl#3=deny any any deernet-iv"
```

This example shows how to apply an output ACL in ASCII format to a port for the duration of this connection:

```
cisco-avpair= "ip:outacl#2=deny ip 10.10.10.10 0.0.255.255 any"
```

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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	radius-server vsa send [accounting authentication]	Enable the switch to recognize and use VSAs as defined by RADIUS IETF attribute 26. • (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. If you enter this command without keywords, both accounting and authentication vendor-specific attributes are used.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your settings.
Step 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 “RADIUS Attributes” appendix in the *Cisco IOS Security Configuration Guide for 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	radius-server host {hostname ip-address} 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.

	Command	Purpose
Step 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your settings.
Step 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.

Controlling Switch Access with Kerberos

This section describes how to enable and configure the Kerberos security system, which authenticates requests for network resources by using a trusted third party. To use this feature, the cryptographic (that is, supports encryption) version of the switch software must be installed on your switch. You must obtain authorization to use this feature and to download the cryptographic software files from Cisco.com. For more information, refer to the release notes for this release.

This section consists of these topics:

- [Understanding Kerberos, page 7-32](#)
- [Kerberos Operation, page 7-34](#)
- [Configuring Kerberos, page 7-35](#)

For Kerberos configuration examples, refer to the “Kerberos Configuration Examples” section in the “Security Server Protocols” chapter of the *Cisco IOS Security Configuration Guide, Release 12.1*, at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/secur_c/scprt2/.

**Note**

For complete syntax and usage information for the commands used in this section, refer to the “Kerberos Commands” section in the “Security Server Protocols” chapter of the *Cisco IOS Security Command Reference, Release 12.1*, at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/secur_r/srprt2/srdkerb.htm.

**Note**

In the Kerberos configuration examples and in the *Cisco IOS Security Command Reference, Release 12.1*, the trusted third party can be a Catalyst 3750 Metro switch that supports Kerberos, that is configured as a network security server, and that can authenticate users by using the Kerberos Protocol.

Understanding Kerberos

Kerberos is a secret-key network authentication protocol, which was developed at the Massachusetts Institute of Technology (MIT). It uses the Data Encryption Standard (DES) cryptographic algorithm for encryption and authentication and authenticates requests for network resources. Kerberos uses the concept of a trusted third party to perform secure verification of users and services. This trusted third party is called the *key distribution center* (KDC).

Kerberos verifies that users are who they claim to be and the network services that they use are what the services claim to be. To do this, a KDC or trusted Kerberos server issues tickets to users. These tickets, which have a limited lifespan, are stored in user credential caches. The Kerberos server uses the tickets instead of usernames and passwords to authenticate users and network services.

**Note**

A Kerberos server can be a Catalyst 3750 Metro switch that is configured as a network security server and that can authenticate users by using the Kerberos Protocol.

The Kerberos credential scheme uses a process called *single logon*. This process authenticates a user once and then allows secure authentication (without encrypting another password) wherever that user credential is accepted.

This software release supports Kerberos 5, which allows organizations that are already using Kerberos 5 to use the same Kerberos authentication database on the KDC that they are already using on their other network hosts (such as UNIX servers and PCs).

In this software release, Kerberos supports these network services:

- Telnet
- rlogin
- rsh (Remote Shell Protocol)

Table 7-2 lists the common Kerberos-related terms and definitions:

Table 7-2 Kerberos Terms

Term	Definition
Authentication	A process by which a user or service identifies itself to another service. For example, a client can authenticate to a switch or a switch can authenticate to another switch.
Authorization	A means by which the switch identifies what privileges the user has in a network or on the switch and what actions the user can perform.
Credential	A general term that refers to authentication tickets, such as ticket granting tickets (TGTs) and service credentials. Kerberos credentials verify the identity of a user or service. If a network service decides to trust the Kerberos server that issued a ticket, it can be used in place of re-entering a username and password. Credentials have a default lifespan of 8 hours.
Instance	An authorization level label for Kerberos principals. Most Kerberos principals are of the form <i>user@REALM</i> (for example, <i>smith@EXAMPLE.COM</i>). A Kerberos principal with a Kerberos instance has the form <i>user/instance@REALM</i> (for example, <i>smith/admin@EXAMPLE.COM</i>). The Kerberos instance can be used to specify the authorization level for the user if authentication is successful. The server of each network service might implement and enforce the authorization mappings of Kerberos instances but is not required to do so. Note The Kerberos principal and instance names <i>must</i> be in all lowercase characters. The Kerberos realm name <i>must</i> be in all uppercase characters.
KDC	Key distribution center (KDC) that consists of a Kerberos server and database program that is running on a network host.
Kerberized	A term that describes applications and services that have been modified to support the Kerberos credential infrastructure.
Kerberos realm	A domain consisting of users, hosts, and network services that are registered to a Kerberos server. The Kerberos server is trusted to verify the identity of a user or network service to another user or network service. Note The Kerberos realm name <i>must</i> be in all uppercase characters.
Kerberos server	A daemon that is running on a network host. Users and network services register their identity with the Kerberos server. Network services query the Kerberos server to authenticate to other network services.
KEYTAB	A password that a network service shares with the KDC. In Kerberos 5 and later Kerberos versions, the network service authenticates an encrypted service credential by using the key table (KEYTAB) to decrypt it. In Kerberos versions earlier than Kerberos 5, KEYTAB is referred to as server table (SRVTAB).
Principal	Also known as a Kerberos identity, this is who you are or what a service is according to the Kerberos server. Note The Kerberos principal name <i>must</i> be in all lowercase characters.
Service credential	A credential for a network service. When issued from the KDC, this credential is encrypted with the password shared by the network service and the KDC. The password is also shared with the user TGT.

Table 7-2 Kerberos Terms (continued)

Term	Definition
SRVTAB	A password that a network service shares with the KDC. In Kerberos 5 or later Kerberos versions, SRVTAB is referred to as KEYTAB.
TGT	Ticket granting ticket that is a credential that the KDC issues to authenticated users. When users receive a TGT, they can authenticate to network services within the Kerberos realm represented by the KDC.

Kerberos Operation

A Kerberos server can be a Catalyst 3750 Metro switch that is configured as a network security server and that can authenticate remote users by using the Kerberos Protocol. Although you can customize Kerberos in a number of ways, remote users attempting to access network services must pass through three layers of security before they can access network services.

To authenticate to network services by using a Catalyst 3750 Metro switch as a Kerberos server, remote users must follow these steps:

1. [Authenticating to a Boundary Switch, page 7-34](#)
2. [Obtaining a TGT from a KDC, page 7-35](#)
3. [Authenticating to Network Services, page 7-35](#)

Authenticating to a Boundary Switch

This section describes the first layer of security through which a remote user must pass. The user must first authenticate to the boundary switch. This process then occurs:

1. The user opens an un-Kerberized Telnet connection to the boundary switch.
2. The switch prompts the user for a username and password.
3. The switch requests a TGT from the KDC for this user.
4. The KDC sends an encrypted TGT that includes the user identity to the switch.
5. The switch attempts to decrypt the TGT by using the password that the user entered.
 - If the decryption is successful, the user is authenticated to the switch.
 - If the decryption is not successful, the user repeats Step 2 either by re-entering the username and password (noting if Caps Lock or Num Lock is on or off) or by entering a different username and password.

A remote user who initiates a un-Kerberized Telnet session and authenticates to a boundary switch is inside the firewall, but the user must still authenticate directly to the KDC before getting access to the network services. The user must authenticate to the KDC because the TGT that the KDC issues is stored on the switch and cannot be used for additional authentication until the user logs on to the switch.

Obtaining a TGT from a KDC

This section describes the second layer of security through which a remote user must pass. The user must now authenticate to a KDC and obtain a TGT from the KDC to access network services.

For instructions about how to authenticate to a KDC, refer to the “Obtaining a TGT from a KDC” section in the “Security Server Protocols” chapter of the *Cisco IOS Security Configuration Guide, Release 12.1*, at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/secur_c/scprt2/scdkerb.htm#xtocid154005.

Authenticating to Network Services

This section describes the third layer of security through which a remote user must pass. The user with a TGT must now authenticate to the network services in a Kerberos realm.

For instructions about how to authenticate to a network service, refer to the “Authenticating to Network Services” section in the “Security Server Protocols” chapter of the *Cisco IOS Security Configuration Guide, Release 12.1*, at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/secur_c/scprt2/scdkerb.htm#xtocid154006.

Configuring Kerberos

So that remote users can authenticate to network services, you must configure the hosts and the KDC in the Kerberos realm to communicate and mutually authenticate users and network services. To do this, you must identify them to each other. You add entries for the hosts to the Kerberos database on the KDC and add KEYTAB files generated by the KDC to all hosts in the Kerberos realm. You also create entries for the users in the KDC database.

When you add or create entries for the hosts and users, follow these guidelines:

- The Kerberos principal name *must* be in all lowercase characters.
- The Kerberos instance name *must* be in all lowercase characters.
- The Kerberos realm name *must* be in all uppercase characters.



Note

A Kerberos server can be a Catalyst 3750 Metro switch that is configured as a network security server and that can authenticate users by using the Kerberos Protocol.

To set up a Kerberos-authenticated server-client system, follow these steps:

- Configure the KDC by using Kerberos commands.
- Configure the switch to use the Kerberos protocol.

For instructions, refer to the “Kerberos Configuration Task List” section in the “Security Server Protocols” chapter of the *Cisco IOS Security Configuration Guide, Release 12.1*, at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/secur_c/scprt2/scdkerb.htm#xtocid154007.

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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa new-model	Enable AAA.
Step 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 ports.
Step 4	aaa authorization exec local	Configure user AAA authorization, check the local database, and allow the user to run an EXEC shell.
Step 5	aaa authorization network local	Configure user AAA authorization for all network-related service requests.
Step 6	username name [privilege level] {password encryption-type password}	Enter the local database, and establish a username-based authentication system. Repeat this command for each user. - 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.
Step 7	end	Return to privileged EXEC mode.
Step 8	show running-config	Verify your entries.
Step 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. The cryptographic (that is, supports encryption) version of the switch software must be installed on your switch. You must obtain authorization to use this feature and to download the cryptographic software files from Cisco.com. For more information, refer to the release notes for this release.

This section contains this configuration information:

- [Understanding SSH, page 7-37](#)
- [Configuring SSH, page 7-37](#)

**Note**

For complete syntax and usage information for the commands used in this section, refer to the “*Secure Shell Commands*” section in the *Cisco IOS Security Command Reference for Release 12.1*.

Understanding SSH

SSH is a protocol that provides a secure, remote connection to a device. 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.

This software release only supports SSH version 1.

**Note**

The SSH feature in this software release does not support IP Security (IPSec).

SSH supports these user authentication methods:

- TACACS+ (For more information, see the “[Controlling Switch Access with TACACS+](#)” section on [page 7-10](#).)
- RADIUS (For more information, see the “[Controlling Switch Access with RADIUS](#)” section on [page 7-17](#).)
- Local authentication and authorization (For more information, see the “[Configuring the Switch for Local Authentication and Authorization](#)” section on [page 7-36](#).)

For more information about SSH, refer to the “*Configuring Secure Shell*” section in the *Cisco IOS Security Configuration Guide for Release 12.1*.

Configuring SSH

Before configuring SSH, download the cryptographic version of the switch software image from Cisco.com. For more information, refer to the release notes for this release.

For information about configuring SSH and displaying SSH settings, refer to the “*Configuring Secure Shell*” section in the *Cisco IOS Security Configuration Guide for Release 12.1*.



Configuring 802.1x Port-Based Authentication

This chapter describes how to configure IEEE 802.1x port-based authentication on the Catalyst 3750 Metro switch. As LANs extend to hotels, airports, and corporate lobbies, creating insecure environments, 802.1x prevents 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, page 8-1](#)
- [Configuring 802.1x Authentication, page 8-9](#)
- [Displaying 802.1x Statistics and Status, page 8-19](#)

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, page 8-2](#)
- [Authentication Initiation and Message Exchange, page 8-3](#)
- [Ports in Authorized and Unauthorized States, page 8-4](#)
- [Supported Topologies, page 8-5](#)
- [Using 802.1x with Port Security, page 8-6](#)
- [Using 802.1x with Voice VLAN Ports, page 8-6](#)

- Using 802.1x with VLAN Assignment, page 8-7
- Using 802.1x with Guest VLAN, page 8-8
- Using 802.1x with Per-User ACLs, page 8-8

Device Roles

With 802.1x port-based authentication, the devices in the network have specific roles as shown in Figure 8-1.

Figure 8-1 802.1x 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 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 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 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 8-4](#).

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 8-4](#).

The specific exchange of EAP frames depends on the authentication method being used. [Figure 8-2](#) shows a message exchange initiated by the client using the One-Time-Password (OTP) authentication method with a RADIUS server.

Figure 8-2 Message Exchange



Ports in Authorized and Unauthorized States

Depending on the switch port state, the switch can grant a client 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, CDP, and STP 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 port.
- **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 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 8-1 on page 8-2](#)), 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 8-3](#) 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 8-3 Wireless LAN Example



Using 802.1x with Port Security

You can configure 802.1x port and port security in either single-host or multiple-hosts mode. (You also must 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 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 the 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 address is aged, its place in the secure host table can be taken by another host.

If the security violation is caused by the first authenticated host, the port becomes error-disabled and immediately shuts down.

The port security violation modes determine the action for security violations. For more information, see the [“Security Violations” section on page 18-8](#).

- When you manually remove an 802.1x client address from the port security table by using the **no switchport port-security mac-address mac-address** interface configuration command, you should re-authenticate the 802.1x client by using the **dot1x re-authenticate interface interface-id** privileged EXEC command.
- When an 802.1x client logs off, the port transitions 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).

For more information about enabling port security on your switch, see the [“Configuring Port Security” section on page 18-6](#).

Using 802.1x with Voice VLAN Ports

A voice VLAN port is a special access port associated with two VLAN identifiers:

- VVID to carry voice traffic to and from the IP phone. The VVID is used to configure the IP phone connected to the port.
- PVID to carry the data traffic to and from the workstation connected to the switch through the IP phone. The PVID is the native VLAN of the port.

Each port that you configure for a voice VLAN is associated with a PVID and a VVID. This configuration allows voice traffic and data traffic to be separated onto different VLANs. The IP phone uses the VVID for its voice traffic regardless of the authorized or unauthorized state of the port. This allows the phone to work independently of 802.1x authentication.

When you enable the single-host mode, multiple IP phones are allowed on the VVID; only one 802.1x client is allowed on the PVID. When you enable the multiple-hosts mode and when an 802.1x user is authenticated on the primary VLAN, additional clients on the voice VLAN are unrestricted after 802.1x authentication succeeds on the primary VLAN.

A voice VLAN port becomes active when there is link, and the device MAC address appears after the first CDP message from the IP phone. Cisco IP phones do not relay CDP messages from other devices. As a result, if several IP phones are connected in series, the switch recognizes only the one directly connected to it. When 802.1x is enabled on a voice VLAN port, the switch drops packets from unrecognized IP phones more than one hop away.

When 802.1x is enabled on a port, you cannot configure a port VLAN that is equal to a voice VLAN.

For more information about voice VLANs, see the [Chapter 12, “Configuring Voice VLAN.”](#)

Using 802.1x with VLAN Assignment

The switch supports 802.1x with 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. You can use this feature to limit network access for certain users.

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 specifying a VLAN for a routed port, a malformed VLAN ID, a nonexistent or internal (routed port) VLAN ID, or an 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 802.1x and port security are enabled on a port, the port is placed in 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 put into 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 by using the **network** keyword to allow port configuration from the RADIUS server.
- 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 7-29](#).

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 when 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 or multiple-hosts mode.

You can configure any active VLAN except an RSPAN VLAN or a voice VLAN as an 802.1x guest VLAN. The guest VLAN feature is not supported on internal VLANs (routed ports) or trunk ports; it is supported only on access ports.

For more information, see the [“Configuring a Guest VLAN” section on page 8-18](#).

Using 802.1x with Per-User ACLs

You can enable per-user access control lists (ACLs) to provide different levels of network access and service to an 802.1x-authenticated user. When the RADIUS server authenticates a user connected to an 802.1x port, it retrieves the ACL attributes based on the user identity and sends them to the switch. The switch applies the attributes to the 802.1x port for the duration of the user session. The switch removes the per-user ACL configuration when the session is over, if authentication fails, or if a link-down condition occurs. The switch does not save RADIUS-specified ACLs in the running configuration. When the port is unauthorized, the switch removes the ACL from the port.

You can configure router ACLs and input port ACLs. However, a port ACL takes precedence over a router ACL. If you apply input port ACL to a port that belongs to a VLAN, the port ACL takes precedence over an input router ACL applied to the VLAN interface. Incoming packets received on the

port to which a port ACL is applied are filtered by the port ACL. Incoming routed packets received on other ports are filtered by the router ACL. Outgoing routed packets are filtered by the router ACL. To avoid configuration conflicts, you should carefully plan the user profiles stored on the RADIUS server.

RADIUS supports per-user attributes, including vendor-specific attributes. These vendor-specific attributes (VSAs) are in octet-string format and are passed to the switch during the authentication process. The VSAs used for per-user ACLs are `inac1#<n>` for the ingress direction and `outac1#<n>` for the egress direction. MAC ACLs are supported only in the ingress direction. The switch supports VSAs only in the ingress direction. It does not support port ACLs in the egress direction on Layer 2 ports. For more information, see [Chapter 25, “Configuring Network Security with ACLs.”](#)

Use only the extended ACL syntax style to define the per-user configuration stored on the RADIUS server. When the definitions are passed from the RADIUS server, they are created by using the extended naming convention. However, if you use the Filter-Id attribute, it can point to a standard ACL.

You can use the Filter-Id attribute to specify an inbound or outbound ACL that is already configured on the switch. The attribute contains the ACL number followed by `.in` for ingress filtering or `.out` for egress filtering. If the RADIUS server does not allow the `.in` or `.out` syntax, the access list is applied to the outbound ACL by default. Because of limited support of Cisco IOS access lists on the switch, the Filter-Id attribute is supported only for IP ACLs numbered 1 to 199 and 1300 to 2699 (IP standard and IP extended ACLs).

Only one 802.1x-authenticated user is supported on a port. If the multiple-hosts mode is enabled on the port, the per-user ACL attribute is disabled for the associated port.

The maximum size of the per-user ACL is 4000 ASCII characters.

For examples of vendor-specific attributes, see the [“Configuring the Switch to Use Vendor-Specific RADIUS Attributes” section on page 7-29](#). For more information about configuring ACLs, see [Chapter 25, “Configuring Network Security with ACLs.”](#)

To configure per-user ACLs, you need to perform these tasks:

- Enable AAA authentication.
- Enable AAA authorization by using the **network** keyword to allow port configuration from the RADIUS server.
- Enable 802.1x.
- Configure the user profile and VSAs on the RADIUS server.
- Configure the 802.1x port for single-host mode.

Configuring 802.1x Authentication

These sections describe how to configure 802.1x port-based authentication on your switch:

- [Default 802.1x Configuration, page 8-10](#)
- [802.1x Configuration Guidelines, page 8-11](#)
- [Configuring 802.1x Authentication, page 8-11](#) (required)
- [Configuring the Switch-to-RADIUS-Server Communication, page 8-13](#) (required)
- [Configuring Periodic Re-Authentication, page 8-14](#) (optional)
- [Manually Re-Authenticating a Client Connected to a Port, page 8-14](#) (optional)
- [Changing the Quiet Period, page 8-15](#) (optional)
- [Changing the Switch-to-Client Retransmission Time, page 8-15](#) (optional)

- [Setting the Switch-to-Client Frame-Retransmission Number, page 8-16](#) (optional)
- [Configuring the Host Mode, page 8-17](#) (optional)
- [Configuring a Guest VLAN, page 8-18](#) (optional)
- [Resetting the 802.1x Configuration to the Default Values, page 8-18](#) (optional)

Default 802.1x Configuration

[Table 8-1](#) shows the default 802.1x configuration.

Table 8-1 *Default 802.1x Configuration*

Feature	Default Setting
Authentication, authorization, and accounting (AAA)	Disabled.
RADIUS server • IP address • UDP authentication port • Key	• None specified. • 1812. • None specified.
Switch 802.1x enable state	Disabled.
Per-port 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 or Layer 3 features are enabled.
- The 802.1x protocol is supported on Layer 2 static-access ports, voice VLAN ports, and Layer 3 routed 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, an error message appears, and 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, an error message appears, and 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—Do not configure a port that is an active member of an EtherChannel as an 802.1x port. If 802.1x is enabled 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 port. However, 802.1x is disabled until the port is removed as a SPAN or RSPAN destination port. You can enable 802.1x on a SPAN or RSPAN source port.
- You can configure any VLAN except an RSPAN VLAN or a voice VLAN as an 802.1x guest VLAN. The guest VLAN feature is not supported on internal VLANs (routed ports) or 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.

Configuring 802.1x Authentication

To configure 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 per-user ACLs or 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa new-model	Enable AAA.
Step 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 ports. Enter at least one of these keywords: • 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.
Step 4	dot1x system-auth-control	Enable 802.1x authentication globally on the switch.
Step 5	aaa authorization network {default} group radius	(Optional) Configure the switch for user RADIUS authorization for all network-related service requests, such as per-user ACLs or VLAN assignment. Note For per-user ACLs, single-host mode must be configured. This setting is the default.
Step 6	interface interface-id	Specify the port connected to the client that is to be enabled for 802.1x authentication, and enter interface configuration mode.
Step 7	dot1x port-control auto	Enable 802.1x authentication on the port. For feature interaction information, see the “802.1x Configuration Guidelines” section on page 8-11.
Step 8	end	Return to privileged EXEC mode.
Step 9	show dot1x	Verify your entries.
Step 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 a port:

```
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 fastethernet1/0/1
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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	radius-server host {hostname ip-address} auth-port port-number key string	Configure the RADIUS server parameters. 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. The range is 0 to 65536. 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 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 7-28.

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 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port to be configured, and enter interface configuration mode.
Step 3	dot1x reauthentication	Enable periodic re-authentication of the client, which is disabled by default.
Step 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.
Step 5	end	Return to privileged EXEC mode.
Step 6	show dot1x interface <i>interface-id</i>	Verify your entries.
Step 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** interface 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 [“Configuring Periodic Re-Authentication” section on page 8-14](#).

This example shows how to manually re-authenticate the client connected to a port:

```
Switch# dot1x re-authenticate interface fastethernet1/0/1
```

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 **dot1x timeout quiet-period** interface configuration command controls the idle period. 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port to be configured, and enter interface configuration mode.
Step 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show dot1x interface <i>interface-id</i>	Verify your entries.
Step 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port to be configured, and enter interface configuration mode.
Step 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.

	Command	Purpose
Step 4	end	Return to privileged EXEC mode.
Step 5	show dot1xinterface <i>interface-id</i>	Verify your entries.
Step 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port to be configured, and enter interface configuration mode.
Step 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show dot1x interface <i>interface-id</i>	Verify your entries.
Step 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 8-3 on page 8-5](#). 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port to which multiple hosts are indirectly attached, and enter interface configuration mode.
Step 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 port.
Step 4	end	Return to privileged EXEC mode.
Step 5	show dot1x interface <i>interface-id</i>	Verify your entries.
Step 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 802.1x on a port and to allow multiple hosts:

```
Switch(config)# interface fastethernet1/0/1
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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port to be configured, and enter interface configuration mode. For the supported port types, see the “802.1x Configuration Guidelines” section on page 8-11 .
Step 3	dot1x guest-vlan <i>vlan-id</i>	Specify an active VLAN as an 802.1x guest VLAN. The range is 1 to 4094. You can configure any active VLAN except an internal VLANs (routed port), an RSPAN VLAN, or a voice VLAN as an 802.1x guest VLAN.
Step 4	end	Return to privileged EXEC mode.
Step 5	show dot1x interface <i>interface-id</i>	Verify your entries.
Step 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 2 as an 802.1x guest VLAN on a port:

```
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# dot1x guest-vlan 2
```

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. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port to be configured, and enter interface configuration mode.
Step 3	dot1x default	Reset the configurable 802.1x parameters to the default values.
Step 4	end	Return to privileged EXEC mode.
Step 5	show dot1x interface <i>interface-id</i>	Verify your entries.
Step 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 ports, use the **show dot1x all statistics** privileged EXEC command. To display 802.1x statistics for a specific port, 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 port, 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.



Configuring Interface Characteristics

This chapter defines the types of interfaces on the Catalyst 3750 Metro switch and describes how to configure them.

The chapter has these sections:

- [Understanding Interface Types, page 9-1](#)
- [Using Interface Configuration Mode, page 9-6](#)
- [Configuring Ethernet Interfaces, page 9-11](#)
- [Configuring Layer 3 Interfaces, page 9-18](#)
- [Configuring the System MTU, page 9-19](#)
- [Monitoring and Maintaining the Interfaces, page 9-20](#)



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 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 physical interface characteristics.

These sections are included:

- [Port-Based VLANs, page 9-2](#)
- [Switch Ports, page 9-2](#)
- [Routed Ports, page 9-4](#)
- [Switch Virtual Interfaces, page 9-4](#)
- [EtherChannel Port Groups, page 9-5](#)
- [Connecting Interfaces, page 9-5](#)

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 10, “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 database 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.
- For a tunnel port, set and define the VLAN ID for the customer-specific VLAN tag. See [Chapter 13, “Configuring IEEE 802.1Q and Layer 2 Protocol Tunneling.”](#)

Switch Ports

Switch ports are Layer 2-only interfaces associated with a physical port. Switch ports belong to one or more VLANs. A switch port can be an access port, a trunk port, or a tunnel 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 switchport mode by negotiating with the port on the other end of the link. You must manually configure tunnel ports as part of an asymmetric link connected to an 802.1Q trunk port. Switch ports are used for managing the physical interface and associated Layer 2 protocols and do not handle routing or bridging.

Configure switch ports by using the **switchport** interface configuration commands. For detailed information about configuring access port and trunk port characteristics, see [Chapter 10, “Configuring VLANs.”](#) For more information about tunnel ports, see [Chapter 13, “Configuring IEEE 802.1Q and Layer 2 Protocol Tunneling.”](#)

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 a tagged packet (Inter-Switch Link [ISL] or 802.1Q tagged), the packet is dropped, and the source address is not learned.

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 VMPS can be a Catalyst 6000 series switch; the Catalyst 3750 Metro switch cannot be a VMPS server.

You can also configure an access port with an attached Cisco IP Phone to use one VLAN for voice traffic and another VLAN for data traffic from a device attached to the phone. For more information about voice VLAN ports, see [Chapter 12, “Configuring Voice VLAN.”](#)

Trunk Ports

A trunk port carries the traffic of multiple VLANs and by default is a member of all VLANs in the VLAN database. Two types of trunk ports are supported:

- In an ISL trunk port, all received packets are expected to be encapsulated with an ISL header, and all transmitted packets are sent with an ISL header. Native (non-tagged) frames received from an ISL trunk port are dropped.
- 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, all possible VLANs (VLAN ID 1 to 4094) are in the allowed list. 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 10, “Configuring VLANs.”](#)

Tunnel Ports

Tunnel ports are used in 802.1Q tunneling to segregate the traffic of customers in a service provider network from other customers who appear to be on the same VLAN. You configure an asymmetric link from a tunnel port on a service provider edge switch to an 802.1Q trunk port on the customer switch. Packets entering the tunnel port on the edge switch, already 802.1Q-tagged with the customer VLANs, are encapsulated with another layer of 802.1Q tag (called the metro tag) containing a VLAN ID unique in the service provider network for each customer. The double-tagged packets go through the service-provider network, keeping the original customer VLANs separate from those of other customers. At the outbound interface, also a tunnel port, the metro tag is removed, and the original VLAN numbers from the customer network are retrieved.

Tunnel ports cannot be trunk ports or access ports and must belong to a VLAN unique for each customer.

For more information about tunnel ports, see [Chapter 13, “Configuring IEEE 802.1Q and Layer 2 Protocol Tunneling.”](#)

Routed Ports

A routed port is a physical port that acts like a port on a router; it does not have to be connected to a router. A routed port is not associated with a particular VLAN, as is an access port. A routed port behaves like a regular router interface, except that it does not support VLAN subinterfaces. Routed ports can be configured with a Layer 3 routing protocol. A routed port is a Layer 3 interface only and does not support Layer 2 protocols, such as DTP and STP.

Configure routed ports by putting the interface into Layer 3 mode with the **no switchport** interface configuration command. Then assign an IP address to the port, enable routing, and assign routing protocol characteristics by using the **ip routing** and **router protocol** global configuration commands.



Caution

Entering a **no switchport** interface configuration command shuts down the interface and then re-enables it, which might generate messages on the device to which the interface is connected.

The number of routed ports that you can configure is not limited by software. However, the interrelationship between this number and the number of other features being configured might impact CPU performance because of hardware limitations. See the [“Configuring Layer 3 Interfaces” section on page 9-18](#) for information about what happens when hardware resource limitations are reached.

For more information about IP unicast and multicast routing and routing protocols, see [Chapter 28, “Configuring IP Unicast Routing,”](#) and [Chapter 31, “Configuring IP Multicast Routing.”](#)

Switch Virtual Interfaces

A switch virtual interface (SVI) represents a VLAN of switch ports as one interface to the routing or bridging function in the system. Only one SVI can be associated with a VLAN, but you need to configure an SVI for a VLAN only when you wish to route between VLANs, to fallback-bridge nonroutable protocols between VLANs, or to provide IP host connectivity to the switch. By default, an SVI is created for the default VLAN (VLAN 1) to permit remote switch administration. Additional SVIs must be explicitly configured. SVIs provide IP host connectivity only to the system; in Layer 3 mode, you can configure routing across SVIs.

Although the switch stack supports a total of 1005 VLANs (and SVIs), the interrelationship between the number of SVIs and routed ports and the number of other features being configured might impact CPU performance because of hardware limitations. See the [“Configuring Layer 3 Interfaces” section on page 9-18](#) for information about what happens when hardware resource limitations are reached.

SVIs are created the first time that you enter the **vlan** interface configuration command for a VLAN interface. The VLAN corresponds to the VLAN tag associated with data frames on an ISL or 802.1Q encapsulated trunk or the VLAN ID configured for an access port. Configure a VLAN interface for each VLAN for which you want to route traffic, and assign it an IP address. For more information, see the [“Manually Assigning IP Information” section on page 3-9](#).



Note

When you create an SVI, it does not become active until it is associated with a physical port.

SVIs support routing protocols and bridging configurations. For more information about configuring IP routing, see [Chapter 28, “Configuring IP Unicast Routing,”](#) [Chapter 31, “Configuring IP Multicast Routing,”](#) and [Chapter 33, “Configuring Fallback Bridging.”](#)

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, group multiple access ports into one logical access port, group multiple tunnel ports into one logical tunnel port, or group multiple routed ports into one logical routed 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), and the Port Aggregation Protocol (PAgP), 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 3 interfaces, you manually create the logical interface by using the **interface port-channel** global configuration command. Then you manually assign an interface to the EtherChannel by using the **channel-group** interface configuration command. For Layer 2 interfaces, use the **channel-group** interface configuration command to dynamically create the port-channel logical interface. This command binds the physical and logical ports together. For more information, see [Chapter 27, “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. With a standard Layer 2 switch, ports in different VLANs have to exchange information through a router. In the configuration shown in [Figure 9-1](#), 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 9-1 Connecting VLANs with Layer 2 Switches



By using the switch with routing enabled (as a Layer 3 switch), when you configure VLAN 20 and VLAN 30 each with an SVI to which an IP address is assigned, packets can be sent from Host A to Host B directly through the switch with no need for an external router ([Figure 9-2](#)).

Figure 9-2 Connecting VLANs with a Layer 3 Switch

The switch supports two methods of forwarding traffic between interfaces: routing and fallback bridging. Whenever possible, to maintain high performance, forwarding is done by the switch hardware. However, only IP version 4 packets with Ethernet II encapsulation can be routed in hardware. Non-IP traffic and traffic with other encapsulation methods can be fallback-bridged by hardware.

- The routing function can be enabled on all SVIs and routed ports. The switch routes only IP traffic. When IP routing protocol parameters and address configuration are added to an SVI or routed port, any IP traffic received from these ports is routed. For more information, see [Chapter 28, “Configuring IP Unicast Routing,”](#) [Chapter 31, “Configuring IP Multicast Routing,”](#) and [Chapter 32, “Configuring MSDP.”](#)
- Fallback bridging forwards traffic that the switch does not route or traffic belonging to a nonroutable protocol, such as DECnet. Fallback bridging connects multiple VLANs into one bridge domain by bridging between two or more SVIs or routed ports. When configuring fallback bridging, you assign SVIs or routed ports to bridge groups with each SVI or routed port assigned to only one bridge group. All interfaces in the same group belong to the same bridge domain. For more information, see [Chapter 33, “Configuring Fallback Bridging.”](#)

Using Interface Configuration Mode

The switch supports these interface types:

- Physical ports—including switch ports and routed ports
- VLANs—switch virtual interfaces
- Port-channels—EtherChannel of interfaces

You can also configure a range of interfaces (see the [“Configuring a Range of Interfaces”](#) section on [page 9-8](#)).

To configure a physical interface (port), enter interface configuration mode, and specify the interface type, switch number, module number, and switch port number.

- Type—Fast Ethernet (fastethernet or fa) for 10/100 Mbps Ethernet or Gigabit Ethernet (gigabitethernet or gi) for small form-factor pluggable (SFP) Gigabit Ethernet interfaces.
- Switch number—For the Catalyst 3750 Metro switch, this number is always 1.

- Module number—The module or slot number on the switch. The module number is zero (0) for Fast Ethernet interfaces and standard SFP module interfaces or 1 for enhanced-services (ES) SFP module interfaces.
- 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, fastethernet 1/0/1, fastethernet 1/0/2. The numbers start again from 1 when the type of interface changes. The standard SFP module interfaces are gigabitethernet 1/0/1 and gigabitethernet 1/0/2, and the ES SFP module interfaces are gigabitethernet 1/1/1 and gigabitethernet 1/1/2.

You can identify physical interfaces by physically checking the interface location on the switch. You can also use the IOS **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.

Procedures for Configuring Interfaces

These general instructions apply to all interface configuration processes.

-
- Step 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)#
```

- Step 2** Enter the **interface** global configuration command. Identify the interface type and number. For physical interfaces, this includes the switch number (1) and the number of the port to be configured:

```
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)#
```



Note You do not need to add a space between the physical interface type and the interface number. For example, in the preceding line, you can specify **gigabitethernet 1/0/1**, **gigabitethernet1/0/1**, **gi 1/0/1**, or **gi1/0/1**.

- Step 3** Follow each **interface** command with the interface configuration commands that the interface requires. The commands that 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.

- Step 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 9-20](#).
-

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.

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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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. - 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 9-9. - In a comma-separated <i>port-range</i>, you must enter the interface type for each entry and enter spaces before and after the comma. - In a hyphen-separated <i>port-range</i>, you do not need to re-enter the interface type, but you must enter a space before the hyphen.
Step 3		You can now use the normal configuration commands to apply the configuration parameters to all interfaces in the range.
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces [<i>interface-id</i>]	Verify the configuration of the interfaces in the range.
Step 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
 - fastethernet** *switch/module/{first port} - {last port}*, where *switch* is **1** and *module* is **0**
 - gigabitethernet** *switch/module/{first port} - {last port}*, where *switch* is **1** and *module* is **0** for the standard SFP module ports and **1** for the ES SFP module ports
 - port-channel** *port-channel-number* - *port-channel-number*, where *port-channel-number* is from 1 to 12



Note When you use the **interface range** command with port channels, the first and last port channel number must be active port channels.

- You must add a space between the first interface number and the hyphen when using the **interface range** command. For example, the command **interface range fastethernet 1/0/1 - 5** is a valid range; the command **interface range fastethernet 1/0/1-5** is not a valid range.

- The **interface range** command only works with VLAN interfaces that have been configured with the **interface vlan** command. The **show running-config** privileged EXEC command displays the configured VLAN interfaces. VLAN interfaces not displayed by the **show running-config** command cannot be used with the **interface range** command.
- All interfaces defined as in a range must be the same type (all Fast Ethernet ports, all Gigabit Ethernet SFP module ports, all EtherChannel ports, or all VLANs), but you can enter multiple ranges in a command.

This example shows how to use the **interface range** global configuration command to set the speed on 10/100 ports 1 to 5 to 100 Mbps:

```
Switch# configure terminal
Switch(config)# interface range fastethernet1/0/1 - 5
Switch(config-if-range)# speed 100
```

This example shows how to use a comma to add different interface type strings to the range to enable Fast Ethernet ports in the range 1 to 3 and Gigabit Ethernet standard SFP module ports 1 and 2 to receive flow control pause frames:

```
Switch# configure terminal
Switch(config)# interface range fastethernet1/0/1 - 3 , gigabitethernet1/0/1 - 2
Switch(config-if-range)# flowcontrol receive on
```

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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	define interface-range <i>macro_name</i> <i>interface-range</i>	Define the interface-range macro, and save it in NVRAM. • The <i>macro_name</i> is a 32-character maximum character string. • A macro can contain up to five comma-separated interface ranges. • Each <i>interface-range</i> must consist of the same port type.
Step 3	interface range macro <i>macro_name</i>	Select the interface range to be configured 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.
Step 4	end	Return to privileged EXEC mode.

	Command	Purpose
Step 5	show running-config include define	Show the defined interface range macro configuration.
Step 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
 - fastethernet** *switch/module/{first port} - {last port}*, where *switch* is **1** and *module* is **0**
 - gigabitethernet** *switch/module/{first port} - {last port}*, where *switch* is **1** and *module* is **0** for the standard SFP module ports and **1** for the ES SFP module ports.
 - port-channel** *port-channel-number* - *port-channel-number*, where *port-channel-number* is from 1 to 12.



Note When you use the interface ranges with port channels, the first and last port channel number must be active port channels.

- You must add a space between the first interface number and the hyphen when entering an *interface-range*. For example, **fastethernet 1/0/1 - 5** is a valid range; **fastethernet 1/0/1-5** is not a valid range.
- The VLAN interfaces must have been configured with the **interface vlan** command. The **show running-config** privileged EXEC command displays the configured VLAN interfaces. VLAN interfaces not displayed by the **show running-config** command cannot be used as *interface-ranges*.
- All interfaces defined as in a range must be the same type (all Fast Ethernet ports, all Gigabit Ethernet SFP module 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 Fast Ethernet ports 1 to 4 and to verify the macro configuration:

```
Switch# configure terminal
Switch(config)# define interface-range enet_list fastethernet1/0/1 - 4
Switch(config)# end
Switch# show running-config | include define
define interface-range enet_list FastEthernet1/0/1 - 4
```

This example shows how to create a multiple-interface macro named *macro1*:

```
Switch# configure terminal
Switch(config)# define interface-range macro1 fastethernet1/0/1 - 2, fastethernet1/0/5 - 7
Switch(config)# end
```

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 was deleted.

```
Switch# configure terminal
Switch(config)# no define interface-range enet_list
Switch(config)# end
Switch# show run | include define
Switch#
```

Configuring Ethernet Interfaces

These sections describe the default interface configuration and the optional features that you can configure on most physical interfaces:

- [Default Ethernet Interface Configuration, page 9-11](#)
- [Configuring Interface Speed and Duplex Mode, page 9-12](#)
- [Configuring IEEE 802.3z Flow Control, page 9-15](#)
- [Configuring Auto-MDIX on a Port, page 9-16](#)
- [Adding a Description for an Interface, page 9-17](#)

Default Ethernet Interface Configuration

[Table 9-1](#) shows the Ethernet interface default configuration, including some features that apply only to Layer 2 interfaces. For more details on the VLAN parameters listed in the table, see [Chapter 10, “Configuring VLANs.”](#) For details on controlling traffic to the port, see [Chapter 18, “Configuring Port-Based Traffic Control.”](#)



Note

To configure Layer 2 parameters, if the interface is in Layer 3 mode, you must enter the **switchport** interface configuration command without any parameters to put the interface into Layer 2 mode. This shuts down the interface and then re-enables it, which might generate messages on the device to which the interface is connected. Furthermore, when you use this command to put the interface into Layer 2 mode, you are deleting any Layer 3 characteristics configured on the interface.

Table 9-1 Default Layer 2 Ethernet Interface Configuration

Feature	Default Setting
Operating mode	Layer 2 or switching mode (switchport command).
Allowed VLAN range	VLANs 1 – 4094.
Default VLAN (for access ports)	VLAN 1 (Layer 2 interfaces only).
Native VLAN (for 802.1Q trunks)	VLAN 1 (Layer 2 interfaces only).
VLAN trunking	Switchport mode dynamic auto (supports DTP) (Layer 2 interfaces only).
Port enable state	All ports are enabled.
Port description	None defined.
Speed	Autonegotiate.
Duplex mode	Autonegotiate.

Table 9-1 Default Layer 2 Ethernet Interface Configuration (continued)

Feature	Default Setting
Flow control	Flow control is set to receive: off . It is always off for sent packets.
EtherChannel (PAgP)	Disabled on all Ethernet ports. See Chapter 27, “Configuring EtherChannels.”
Port blocking (unknown multicast and unknown unicast traffic)	Disabled (not blocked) (Layer 2 interfaces only). See the “Configuring Port Blocking” section on page 18-5.
Broadcast, multicast, and unicast storm control	Disabled. See the “Default Storm Control Configuration” section on page 18-3.
Protected port	Disabled (Layer 2 interfaces only). See the “Configuring Protected Ports” section on page 18-4.
Port security	Disabled (Layer 2 interfaces only). See the “Default Port Security Configuration” section on page 18-9. L2
Port Fast	Disabled.
Automatic medium-dependent-interface crossover (Auto-MDIX)	Disabled.

Configuring Interface Speed and Duplex Mode

Ethernet interfaces on the switch operate at 10, 100, or 1000 Mbps and in either full- or half-duplex mode. In full-duplex mode, two stations can send and receive traffic at the same time. Normally, 10-Mbps ports operate in half-duplex mode, which means that stations can either receive or send traffic.

Switch models include combinations of Fast Ethernet (10/100-Mbps) ports or Gigabit Ethernet SFP standard and ES module slots supporting Gigabit SFP modules.

- You can configure interface speed on Fast Ethernet (10/100-Mbps) ports. You can configure duplex mode to full, half, or autonegotiate on Fast Ethernet ports.
- You cannot configure speed or duplex mode on SFP ports, but you can configure speed to not negotiate (**nonegotiate**) if connected to a device that does not support autonegotiation. However, when a 1000BASE-T SFP module is in the SFP module port, you can configure speed as 10, 100, or 1000 Mbps or to autonegotiate, and you can configure duplex mode to auto or full.



Note

ES ports do not support 10 or 100 Mbps operation. When a 1000BASE-T SFP module is inserted in an ES module port, the speed defaults to 1000. Setting the speed to **auto** can result in severe data loss if the interface is connected to a 10/100 interface on another switch.

These sections describe how to configure the interface speed and duplex mode:

- [Configuration Guidelines, page 9-13](#)
- [Setting the Interface Speed and Duplex Parameters, page 9-13](#)

Configuration Guidelines

When configuring an interface speed and duplex mode, note these guidelines:

- If both ends of the line support autonegotiation, we highly recommend the default setting of **auto** negotiation.
- If one interface supports autonegotiation and the other end does not, configure duplex and speed on both interfaces; do not use the **auto** setting on the supported side.
- For 10/100 Mbps ports, if both speed and duplex are set to specific values, the link operates at the negotiated speed and duplex value.
- You cannot configure duplex mode on SFP module ports; they operate only in full-duplex mode. However, when a 1000BASE-T SFP module is inserted in an SFP module port, you can configure the duplex mode to **full** or **auto** and half-duplex mode is supported with the auto configuration.
- You cannot configure speed on SFP module ports, except to **nonegotiate**. However, when a 1000BASE-T SFP module is in the SFP module port, the speed can be configured to **10**, **100**, **1000**, or **auto**, but not **nonegotiate**.



Caution

The ES ports do not support 10 or 100 Mbps operation. Setting the speed on an ES port to 10 or 100 Mbps, or to **auto** if the interface is connected to a 10/100 interface on another switch, can result in severe data loss if the interface is connected to a 10/100 interface on another switch.

- When STP is enabled and a port is reconfigured, the switch can take up to 30 seconds to check for loops. The port LED is amber while STP reconfigures.



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

Beginning in privileged EXEC mode, follow these steps to set the speed and duplex mode for a physical interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode and the physical interface identification.

	Command	Purpose
Step 3	speed { 10 100 1000 auto nonegotiate }	Enter the appropriate speed parameter for the port: - Enter 10, 100, or 1000 to set a specific speed for the port. The 1000 keyword is not available on 10/100 Mbps ports. These keywords are available on SFP module ports only when a 1000BASE-T SFP module is inserted. - Enter auto to enable the port to autonegotiate speed with the device connected to the port. This option is available on SFP module ports only when a 1000BASE-T SFP module is inserted. - The nonegotiate keyword is available only for SFP module ports. SFP module ports operate only at 1000 Mbps but can be configured to not negotiate if connected to a device that does not support autonegotiation. Note When a 1000BASE-T SFP module is in the SFP module port, the speed can be configured to 10, 100, 1000, or auto, but not nonegotiate. However, ES ports operate only at 1000 Mbps. If you configure a 1000BASE-T SFP module in an ES port to operate at 10 or 100 Mbps or to autonegotiate to 10 or 100 Mbps, the interface experiences large numbers of traffic errors.
Step 4	duplex { auto full half }	Enter the duplex parameter for the interface. This command is not available on SFP module ports unless a 1000BASE-T SFP module is inserted, when the mode can then be configured to auto or full.
Step 5	end	Return to privileged EXEC mode.
Step 6	show interfaces <i>interface-id</i>	Display the interface speed and duplex mode configuration.
Step 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 speed to 10 Mbps and the duplex mode to half on an port:

```
Switch# configure terminal
Switch(config)# interface fasttetherenet1/0/3
Switch(config-if)# speed 10
Switch(config-if)# duplex half
```

Configuring IEEE 802.3z Flow Control

Flow control enables connected 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 by sending a pause frame. Upon receipt of a pause frame, the sending device stops sending any data packets, which prevents any loss of data packets during the congestion period.



Note

Catalyst 3750 Metro switch ports are capable of receiving, but not sending, pause frames.

You use the **flowcontrol** interface configuration command to set the interface's ability to **receive** pause frames to **on**, **off**, or **desired**. The default state is **off**.

When set to **desired**, an interface can operate with an attached device that is required to send flow-control packets or with an attached device that is not required to but can send flow-control packets.

These rules apply to flow control settings on the device:

- **receive on (or desired):** 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:** 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 on the command settings and the resulting flow control resolution on local and remote ports, refer to the **flowcontrol** interface configuration command in the command reference for this release.

Beginning in privileged EXEC mode, follow these steps to configure flow control on an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode
Step 2	interface <i>interface-id</i>	Enter interface configuration mode and the physical interface to be configured.
Step 3	flowcontrol {receive} {on off desired}	Configure the flow control mode for the port.
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces <i>interface-id</i>	Verify the interface flow control settings.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable flow control, use the **flowcontrol receive off** interface configuration command.

This example shows how to turn on flow control on a port:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# flowcontrol receive on
Switch(config-if)# end
```

Configuring Auto-MDIX on a Port

When automatic medium-dependent-interface crossover (Auto-MDIX) is enabled on an port, the port automatically detects the required cable connection type (straight through or crossover) and configures the connection appropriately. When connecting switches without the Auto-MDIX feature, you must use straight-through cables to connect to devices such as servers, workstations, or routers and crossover cables to connect to other switches or repeaters. With Auto-MDIX enabled, you can use either type of cable to connect to other devices, and the interface automatically corrects for any incorrect cabling. For more information about cabling requirements, refer to the hardware installation guide.

Auto-MDIX is disabled by default. When you enable Auto-MDIX, you must also set the speed and duplex on the port to **auto** in order for the feature to operate correctly. Auto-MDIX is supported on all 10/100 Mbps ports and on 10/100/1000 BASE-T/TX SFP module ports. It is not supported on 1000 BASE-SX or -LX SFP module ports.



Caution

ES ports do not support 10 or 100 Mbps operation. When a 1000BASE-T SFP module is inserted in an ES module port, if you set the speed to **auto** to enable Auto-MDIX, be sure that the interface is connected to a Gigabit interface. Setting the speed to **auto** can result in severe data loss if the interface is connected to a 10/100 interface on another switch.

Table 9-2 shows the link states that results from auto-MDIX settings and correct and incorrect cabling.

Table 9-2 Link Conditions and Auto-MDIX Settings

Local Side Auto-MDIX	Remote Side Auto-MDIX	With Correct Cabling	With Incorrect Cabling
On	On	Link up	Link up
On	Off	Link up	Link up
Off	On	Link up	Link up
Off	Off	Link up	Link down

Beginning in privileged EXEC mode, follow these steps to configure Auto-MDIX on a port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode
Step 2	interface <i>interface-id</i>	Enter interface configuration mode for the physical interface to be configured.
Step 3	speed auto	Configure the port to autonegotiate speed with the connected device.
Step 4	duplex auto	Configure the port to autonegotiate duplex mode with the connected device.
Step 5	mdix auto	Enable Auto-MDIX on the port. Note On an ES port, setting the speed to auto can result in severe data loss if the port is connected to a 10/100 interface on another switch.
Step 6	end	Return to privileged EXEC mode.
Step 7	show controllers ethernet-controller <i>interface-id</i> phy	Verify the operational state of the Auto-MDIX feature on the interface.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable Auto-MDIX, use the **no mdix auto** interface configuration command.

This example shows how to enable Auto-MDIX on a port:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# speed auto
Switch(config-if)# duplex auto
Switch(config-if)# mdix auto
Switch(config-if)# end
```

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 privileged EXEC commands: **show configuration**, **show running-config**, and **show interfaces**.

Beginning in privileged EXEC mode, follow these steps to add a description for an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the interface for which you are adding a description.
Step 3	description <i>string</i>	Add a description (up to 240 characters) for an interface.
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces <i>interface-id</i> description or show running-config	Verify your entry.
Step 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 a port and to verify the description:

```
Switch# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# description Connects to Marketing
Switch(config-if)# end
Switch# show interfaces gigabitethernet1/0/2 description
Interface Status      Protocol Description
Gi1/0/2 admin down    down    Connects to Marketing
```

Configuring Layer 3 Interfaces

The switch supports these types of Layer 3 interfaces:

- SVIs: You should configure SVIs for any VLANs for which you want to route traffic. SVIs are created when you enter a VLAN ID following the **interface vlan** global configuration command. To delete an SVI, use the **no interface vlan** global configuration command.

**Note**

When you create an SVI, it does not become active until it is associated with a physical port. For information about assigning Layer 2 ports to VLANs, see [Chapter 10, “Configuring VLANs.”](#)

- Routed ports: Routed ports are physical ports configured to be in Layer 3 mode by using the **no switchport** interface configuration command.
- Layer 3 EtherChannel ports: EtherChannel interfaces made up of routed ports.

EtherChannel port interfaces are described in [Chapter 27, “Configuring EtherChannels.”](#)

A Layer 3 switch can have an IP address assigned to each routed port and SVI.

There is no defined limit to the number of SVIs and routed ports that can be configured in a switch. However, the interrelationship between the number of SVIs and routed ports and the number of other features being configured might have an impact on CPU usage because of hardware limitations. If the switch is using maximum hardware resources, attempts to create a routed port or SVI have these results:

- If you try to create a new routed port, the switch generates a message that there are not enough resources to convert the interface to a routed port, and the interface remains as a switchport.
- If you try to create an extended-range VLAN, an error message is generated, and the extended-range VLAN is rejected.
- If the switch is notified by VLAN Trunking Protocol (VTP) of a new VLAN, it sends a message that there are not enough hardware resources available and shuts down the VLAN. The output of the **show vlan** user EXEC command shows the VLAN in a suspended state.
- If the switch attempts to boot up with a configuration that has more VLANs and routed ports than hardware can support, the VLANs are created, but the routed ports are shut down, and the switch sends a message that this was due to insufficient hardware resources.

All Layer 3 interfaces require an IP address to route traffic. This procedure shows how to configure an interface as a Layer 3 interface and how to assign an IP address to an interface.

**Note**

If the physical port is in Layer 2 mode (the default), you must enter the **no switchport** interface configuration command to put the interface into Layer 3 mode. Entering a **no switchport** command disables and then re-enables the interface, which might generate messages on the device to which the interface is connected.

Beginning in privileged EXEC mode, follow these steps to configure a Layer 3 interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface {{ fastethernet gigabitethernet } <i>interface-id</i> } { vlan <i>vlan-id</i> } { port-channel <i>port-channel-number</i> }	Enter interface configuration mode, and enter the interface to be configured as a Layer 3 interface.
Step 3	no switchport	For physical ports only, enter Layer 3 mode.
Step 4	ip address <i>ip_address subnet_mask</i>	Configure the IP address and IP subnet.
Step 5	no shutdown	Enable the interface.
Step 6	end	Return to privileged EXEC mode.
Step 7	show interfaces [<i>interface-id</i>] show ip interface [<i>interface-id</i>] show running-config interface [<i>interface-id</i>]	Verify the configuration.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove an IP address from an interface, use the **no ip address** interface configuration command.

This example shows how to configure a port as a routed port and to assign it an IP address:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# no switchport
Switch(config-if)# ip address 192.20.135.21 255.255.255.0
Switch(config-if)# no shutdown
```

Configuring the System MTU

The default maximum transmission unit (MTU) size for frames received and transmitted on all interfaces on the switch is 1500 bytes. You can increase the MTU size for all interfaces operating at 10 or 100 Mbps by using the **system mtu** global configuration command. You can increase the MTU size to support jumbo frames on all Gigabit Ethernet interfaces by using the **system mtu jumbo** global configuration command. Gigabit Ethernet ports are not affected by the **system mtu** command; 10/100 ports are not affected by the **system jumbo mtu** command.

The switch supports a maximum MTU size for 10/100 interfaces of 1546 bytes. The maximum MTU size for Gigabit interfaces is 9000 bytes.

You cannot set the MTU size for an individual interface; you set it for all Fast Ethernet or all Gigabit Ethernet interfaces on the switch. When you change the MTU size, you must reset the switch before the new configuration takes effect.

The size of frames that can be received by the switch CPU is limited to 1500 bytes, no matter what value was entered with the **system mtu** or **system mtu jumbo** commands. Although frames that are forwarded or routed typically are not received by the CPU, in some cases packets are sent to the CPU, such as traffic sent to control traffic, SNMP, or Telnet, or routing protocols.



Note

If Gigabit Ethernet interfaces are configured to accept frames greater than the 10/100 interfaces, jumbo frames ingressing on a Gigabit Ethernet interface and egressing on a 10/100 interface are dropped.

Beginning in privileged EXEC mode, follow these steps to change MTU size for all 10/100 or Gigabit Ethernet interfaces:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	system mtu bytes	(Optional) Change the MTU size for all interfaces on the switch that are operating at 10 or 100 Mbps. The range is from 1500 to 1546 bytes; the default is 1500 bytes.
Step 3	system mtu jumbo bytes	(Optional) Change the MTU size for all Gigabit Ethernet interfaces on the switch. The range is from 1500 to 9000 bytes; the default is 1500 bytes.
Step 4	end	Return to privileged EXEC mode.
Step 5	copy running-config startup-config	Save your entries in the configuration file.
Step 6	reload	Reload the operating system.

If you enter a value that is outside the allowed range for the specific type of interface, the value is not accepted.

Once the switch reloads, you can verify your settings by entering the **show system mtu** privileged EXEC command.

This example shows how to set the maximum packet size for Gigabit Ethernet interfaces to 1800 bytes:

```
Switch(config)# system jumbo mtu 1800
Switch(config)# exit
Switch# reload
```

This example shows the response when you try to set Gigabit Ethernet interfaces to an out-of-range number:

```
Switch(config)# system mtu jumbo 2500
                        ^
% Invalid input detected at '^' marker.
```

Monitoring and Maintaining the Interfaces

You can perform the tasks in these sections to monitor and maintain interfaces:

- [Monitoring Interface Status, page 9-20](#)
- [Clearing and Resetting Interfaces and Counters, page 9-21](#)
- [Shutting Down and Restarting the Interface, page 9-22](#)

Monitoring Interface Status

Commands entered at the privileged EXEC prompt display information about the interface, including the versions of the software and the hardware, the configuration, and statistics about the interfaces. [Table 9-3](#) 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 IOS Interface Command Reference for Release 12.1*.

Table 9-3 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> status [err-disabled]	Display interface status or a list of interfaces in an error-disabled state.
show interfaces [<i>interface-id</i>] switchport	Display administrative and operational status of switching (nonrouting) ports. You can use this command to determine if a port is in routing or switching mode.
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 routing 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.
show controllers ethernet-controller <i>interface-id</i> phy	Verify the operational state of the Auto-MDIX feature on the interface.

Clearing and Resetting Interfaces and Counters

Table 9-4 lists the privileged EXEC mode **clear** commands that you can use to clear counters and reset interfaces.

Table 9-4 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 interface** privileged EXEC command.

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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface { vlan <i>vlan-id</i> } {{ fastethernet gigabitethernet } <i>interface-id</i> } {{ port-channel <i>port-channel-number</i> }	Select the interface to be configured.
Step 3	shutdown	Shut down an interface.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entry.

Use the **no shutdown** interface configuration command to restart the interface.

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 interface** command display.



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 the Catalyst 3750 Metro switch. It includes information about VLAN membership modes, VLAN configuration modes, VLAN trunks, and dynamic VLAN assignment from a 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, page 10-1](#)
- [Configuring Normal-Range VLANs, page 10-4](#)
- [Configuring Extended-Range VLANs, page 10-12](#)
- [Displaying VLANs, page 10-16](#)
- [Configuring VLAN Trunks, page 10-16](#)
- [Configuring VMPS, page 10-28](#)

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 a switch supporting fallback bridging, as shown in [Figure 10-1](#). Because a VLAN is considered a separate logical network, it contains its own bridge MIB information and can support its own implementation of spanning tree. See [Chapter 14, “Configuring STP.”](#)




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 11, “Configuring VTP.”](#)

Figure 10-1 shows an example of VLANs segmented into logically defined networks.

Figure 10-1 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.

Traffic between VLANs must be routed or fallback bridged. The switch can route traffic between VLANs by using switch virtual interfaces (SVIs). An SVI must be explicitly configured and assigned an IP address to route traffic between VLANs. For more information, see the [“Switch Virtual Interfaces”](#) section on page 9-4 and the [“Configuring Layer 3 Interfaces”](#) section on page 9-18.



Note

If you plan to configure many VLANs on the switch and to not enable routing, you can use the **sdm prefer vlan** global configuration command to set the Switch Database Management (SDM) feature to the VLAN template, which configures system resources to support the maximum number of unicast MAC addresses. For more information on the SDM templates, see [Chapter 6, “Configuring SDM Templates,”](#) or refer to the **sdm prefer** command in the command reference for this release.

Supported VLANs

The switch supports 1005 VLANs in VTP client, server, and transparent modes. 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.

Although the switch supports a total of 1005 (normal-range and extended-range) VLANs, the number of routed ports, SVIs, and other configured features affects the use of the switch hardware.

The switch supports per-VLAN spanning-tree plus (PVST+) or rapid PVST+ with a maximum of 128 spanning-tree instances. One spanning-tree instance is allowed per VLAN. See the [“Normal-Range VLAN Configuration Guidelines” section on page 10-6](#) for more information about the number of spanning-tree instances and the number of VLANs. The switch supports both Inter-Switch Link (ISL) and IEEE 802.1Q trunking methods 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 10-1](#) lists the membership modes and membership and VTP characteristics.

Table 10-1 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 10-11 .	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.
Trunk (ISL or IEEE 802.1Q)	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 10-20 .	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.

Table 10-1 Port Membership Modes (continued)

Membership Mode	VLAN Membership Characteristics	VTP Characteristics
Dynamic access	A dynamic-access port can belong to one VLAN (VLAN ID 1 to 4094) and is dynamically assigned by a VMPS. The VMPS can be a Catalyst 5000 or Catalyst 6500 series switch, for example, but never a Catalyst 3750 Metro switch. The Catalyst 3750 Metro switch is a VMPS client. 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 or hub and not to another switch. For configuration information, see the “Configuring Dynamic-Access Ports on VMPS Clients” section on page 10-31.	VTP is required. Configure the VMPS and the client with the same VTP domain name. To participate in VTP, there must be at least one trunk port on the switch connected to a trunk port of a second switch.
Voice VLAN	A voice VLAN port is an access port attached to a Cisco IP Phone, configured to use one VLAN for voice traffic and another VLAN for data traffic from a device attached to the phone. For more information about voice VLAN ports, see Chapter 12, “Configuring Voice VLAN.”	VTP is not required; it has no affect on voice VLAN.
Tunnel (dot1q-tunnel)	Tunnel ports are used for 802.1Q tunneling to maintain customer VLAN integrity across a service provider network. You configure a tunnel port on an edge switch in the service provider network and connect it to an 802.1Q trunk port on a customer interface, creating an asymmetric link. A tunnel port belongs to a single VLAN that is dedicated to tunneling. For more information about tunnel ports, see Chapter 13, “Configuring IEEE 802.1Q and Layer 2 Protocol Tunneling.”	VTP is not required. You manually assign the tunnel port to a VLAN by using the switchport access vlan interface configuration command.

For more detailed definitions of the modes and their functions, see [Table 10-4 on page 10-18](#).

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 5-20.

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 10-12.

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 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 11, “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, page 10-6](#)
- [Normal-Range VLAN Configuration Guidelines, page 10-6](#)
- [VLAN Configuration Mode Options, page 10-7](#)
- [Saving VLAN Configuration, page 10-7](#)
- [Default Ethernet VLAN Configuration, page 10-8](#)
- [Creating or Modifying an Ethernet VLAN, page 10-9](#)
- [Deleting a VLAN, page 10-11](#)
- [Assigning Static-Access Ports to a VLAN, page 10-11](#)

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:

- The switch supports 1005 VLANs in VTP client, server, and transparent modes.
- 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 10-12](#).
- 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 128 spanning-tree instances. If a switch has more active VLANs than supported spanning-tree instances, spanning tree can be enabled on 128 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 15, “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, page 10-7](#)

You access config-vlan mode by entering the **vlan** *vlan-id* global configuration command.

- [VLAN Configuration in VLAN Database Configuration Mode, page 10-7](#)

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 enter an existing VLAN ID to modify the VLAN. You can use the default VLAN configuration ([Table 10-2](#)) 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 10-12](#).

VLAN Configuration in VLAN Database Configuration Mode

To access VLAN database configuration mode, enter the **vlan database** privileged EXEC command. Then enter the **vlan** command with a new VLAN ID to create a VLAN, or enter an existing VLAN ID to modify the VLAN. You can use the default VLAN configuration ([Table 10-2](#)) or enter multiple commands to configure the VLAN. For more information about keywords available in this mode, refer to the **vlan** VLAN database 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** 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 in these ways:

- 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 10-2 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 10-2 Ethernet VLAN Defaults and Ranges

Parameter	Default	Range
VLAN ID	1	1 to 4094. Note Extended-range VLANs (VLAN IDs 1006 to 4094) are not saved in the VLAN database.
VLAN name	VLANxxxx, where xxxx represents four numeric digits (including leading zeros) equal to the VLAN ID number	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	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 10-12.

For the list of default parameters that are assigned when you add a VLAN, see the [“Configuring Normal-Range VLANs”](#) section on page 10-4.

Beginning in privileged EXEC mode, follow these steps to use config-vlan mode to create or modify an Ethernet VLAN:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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 10-12.
Step 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.
Step 4	mtu <i>mtu-size</i>	(Optional) Change the MTU size (or other VLAN characteristic).
Step 5	remote-span	(Optional) Configure the VLAN as the RSPAN VLAN for a remote SPAN session. For more information on remote SPAN, see Chapter 21 , “Configuring SPAN and RSPAN.”
Step 6	end	Return to privileged EXEC mode.
Step 7	show vlan { name <i>vlan-name</i> id <i>vlan-id</i> }	Verify your entries.
Step 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 name**, **no 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
```

You can also create or modify Ethernet VLANs by using the VLAN database configuration mode.

**Note**

VLAN database configuration mode does not support RSPAN VLAN configuration or extended-range VLANs.

Beginning in privileged EXEC mode, follow these steps to use VLAN database configuration mode to create or modify an Ethernet VLAN:

	Command	Purpose
Step 1	vlan database	Enter VLAN database configuration mode.
Step 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. You can create or modify a range of consecutive VLANs by entering vlan <i>first-vlan-id</i> end <i>last-vlan-id</i> . Note When entering a VLAN ID in VLAN database configuration mode, 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.
Step 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.
Step 4	exit	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 5	show vlan { name <i>vlan-name</i> id <i>vlan-id</i> }	Verify your entries.
Step 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.

To return the VLAN name to the default settings, use the **no vlan** *vlan-id* **name** or **no vlan** *vlan-id* **mtu** VLAN database configuration command.

This example shows how to use VLAN 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...
```


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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no vlan <i>vlan-id</i>	Remove the VLAN by entering the VLAN ID.
Step 3	end	Return to privileged EXEC mode.
Step 4	show vlan brief	Verify the VLAN removal.
Step 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 by using VLAN database configuration mode, use the **vlan database** privileged EXEC command to enter VLAN database configuration mode and the **no vlan** *vlan-id* VLAN database 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 cluster 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 10-9.)

Beginning in privileged EXEC mode, follow these steps to assign a port to a VLAN in the VLAN database:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode
Step 2	interface <i>interface-id</i>	Enter the interface to be added to the VLAN.

	Command	Purpose
Step 3	switchport mode access	Define the VLAN membership mode for the port (Layer 2 access port).
Step 4	switchport access vlan <i>vlan-id</i>	Assign the port to a VLAN. Valid VLAN IDs are 1 to 4094.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config interface <i>interface-id</i>	Verify the VLAN membership mode of the interface.
Step 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.
Step 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 a port as an access port in VLAN 2:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan 2
Switch(config-if)# end
```

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 10-3](#) for the actual number of VLANs supported.

This section includes this information about extended-range VLANs:

- [Default VLAN Configuration, page 10-13](#)
- [Extended-Range VLAN Configuration Guidelines, page 10-13](#)
- [Creating an Extended-Range VLAN, page 10-14](#)
- [Creating an Extended-Range VLAN with an Internal VLAN ID, page 10-15](#)

Default VLAN Configuration

See [Table 10-2 on page 10-8](#) for the default configuration for Ethernet VLANs. You can change only the MTU size and remote SPAN configuration state 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 11-12](#). You should 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.
- 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 (128) 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 15, “Configuring MSTP.”](#)
- Each routed port on the switch creates an internal VLAN for its use. These internal VLANs use extended-range VLAN numbers, and the internal VLAN ID cannot be used for an extended-range VLAN. If you try to create an extended-range VLAN with a VLAN ID that is already allocated as an internal VLAN, an error message is generated, and the command is rejected.
 - Because internal VLAN IDs are in the lower part of the extended range, we recommend that you create extended-range VLANs beginning from the highest number (4094) and moving to the lowest (1006) to reduce the possibility of using an internal VLAN ID.
 - Before configuring extended-range VLANs, enter the **show vlan internal usage** privileged EXEC command to see which VLANs have been allocated as internal VLANs.
 - If necessary, you can shut down the routed port assigned to the internal VLAN, which frees up the internal VLAN, and then create the extended-range VLAN and re-enable the port, which then uses another VLAN as its internal VLAN. See the [“Creating an Extended-Range VLAN with an Internal VLAN ID” section on page 10-15](#).
- Although the switch supports a total of 1005 (normal-range and extended-range) VLANs, the number of routed ports, SVIs, and other configured features affects the use of the switch hardware. If you try to create an extended-range VLAN and there are not enough hardware resources available, an error message is generated, and the extended-range VLAN is rejected.

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 10-2](#)) and the MTU size and RSPAN configuration are the only parameters 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 [“Creating an Extended-Range VLAN with an Internal VLAN ID” section on page 10-15](#) before creating the extended-range VLAN.

Beginning in privileged EXEC mode, follow these steps to create an extended-range VLAN:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vtp mode transparent	Configure the switch for VTP transparent mode, disabling VTP.
Step 3	vlan <i>vlan-id</i>	Enter an extended-range VLAN ID and enter config-vlan mode. The range is 1006 to 4094.
Step 4	mtu <i>mtu-size</i>	(Optional) Modify the VLAN by changing the MTU size. Note Although all VLAN commands appear in the CLI help in config-vlan mode, only the mtu <i>mtu-size</i> and remote-span commands are supported for extended-range VLANs.
Step 5	remote-span	(Optional) Configure the VLAN as the RSPAN VLAN. See the “Configuring a VLAN as an RSPAN VLAN” section on page 21-17 .
Step 6	end	Return to privileged EXEC mode.
Step 7	show vlan id <i>vlan-id</i>	Verify that the VLAN has been created.
Step 8	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 10-11](#).

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
```

Creating an Extended-Range VLAN with an Internal VLAN ID

If you enter an extended-range VLAN ID that is already assigned to an internal VLAN, an error message is generated, and the extended-range VLAN is rejected. To manually free an internal VLAN ID, you must temporarily shut down the routed port that is using the internal VLAN ID.

Beginning in privileged EXEC mode, follow these steps to release a VLAN ID that is assigned to an internal VLAN and to create an extended-range VLAN with that ID:

	Command	Purpose
Step 1	show vlan internal usage	Display the VLAN IDs being used internally by the switch. If the VLAN ID that you want to use is an internal VLAN, the display shows the routed port that is using the VLAN ID. Enter that port number in Step 3.
Step 2	configure terminal	Enter global configuration mode.
Step 3	interface <i>interface-id</i>	Enter the interface ID for the routed port that is using the VLAN ID.
Step 4	shutdown	Shut down the port to free the internal VLAN ID.
Step 5	exit	Return to global configuration mode.
Step 6	vtp mode transparent	Set the VTP mode to transparent for creating extended-range VLANs.
Step 7	vlan <i>vlan-id</i>	Enter the new extended-range VLAN ID, and enter config-vlan mode.
Step 8	exit	Exit from config-vlan mode, and return to global configuration mode.
Step 9	interface <i>interface-id</i>	Enter the interface ID for the routed port that you shut down in Step 4.
Step 10	no shutdown	Re-enable the routed port. It will be assigned a new internal VLAN ID.
Step 11	end	Return to privileged EXEC mode.
Step 12	copy running-config startup config	Save your entries in the switch startup configuration file. To save an extended-range VLAN configuration, 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.

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 database configuration command (accessed by entering the **vlan database** privileged EXEC command). For a list of the VLAN IDs on the switch, use the **show vlan** privileged EXEC command.

Table 10-3 lists the commands for monitoring VLANs.

Table 10-3 VLAN Monitoring Commands

Command	Command Mode	Purpose
show	VLAN database configuration	Display status of VLANs in the VLAN database.
show current [vlan-id]	VLAN database configuration	Display status of all or the specified VLAN in the VLAN database.
show interfaces [vlan vlan-id]	Privileged EXEC	Display characteristics for all interfaces or for the specified VLAN configured on the switch.
show vlan [id vlan-id]	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, page 10-16](#)
- [Encapsulation Types, page 10-18](#)
- [Default Layer 2 Ethernet Interface VLAN Configuration, page 10-19](#)
- [Configuring an Ethernet Interface as a Trunk Port, page 10-20](#)
- [Configuring Trunk Ports for Load Sharing, page 10-24](#)


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. Ethernet trunks carry the traffic of multiple VLANs over a single link, and you can extend the VLANs across an entire network.

Two trunking encapsulations are available on all Ethernet interfaces:

- Inter-Switch Link (ISL)—ISL is Cisco-proprietary trunking encapsulation.
- 802.1Q—802.1Q is industry-standard trunking encapsulation.

Figure 10-2 shows a network of switches that are connected by ISL trunks.

Figure 10-2 Switches in an ISL Trunking Environment

You can configure a trunk on a single Ethernet interface or on an EtherChannel bundle. For more information about EtherChannel, see [Chapter 27, “Configuring EtherChannels.”](#)

Ethernet trunk interfaces support different trunking modes (see [Table 10-4](#)). You can set an interface as trunking or nontrunking or to negotiate trunking with the neighboring interface. To autonegotiate trunking, the interfaces must be in the same VTP domain.

Trunk negotiation is managed by the Dynamic Trunking Protocol (DTP), which is a Point-to-Point Protocol. However, some internetworking devices might forward DTP frames improperly, which could cause misconfigurations.

To avoid this, you should configure interfaces connected to devices that do not support DTP to not forward DTP frames, that is, to turn off DTP.

- If you do not intend to trunk across those links, use the **switchport mode access** interface configuration command to disable trunking.
- To enable trunking to a device that does not support DTP, use the **switchport mode trunk** and **switchport nonegotiate** interface configuration commands to cause the interface to become a trunk but to not generate DTP frames. Use the **switchport trunk encapsulation isl** or **switchport trunk encapsulation dot1q** interface to select the encapsulation type on the trunk port.

You can also specify on DTP interfaces whether the trunk uses ISL or 802.1Q encapsulation or if the encapsulation type is autonegotiated. The DTP supports autonegotiation of both ISL and 802.1Q trunks.

**Note**

Tunnel ports do not support DTP. See [Chapter 13, “Configuring IEEE 802.1Q and Layer 2 Protocol Tunneling,”](#) for more information on tunnel ports.

Table 10-4 Layer 2 Interface Modes

Mode	Function
switchport mode access	Puts the interface (access port) into permanent nontrunking mode and negotiates to convert the link into a nontrunk link. The interface becomes a nontrunk interface regardless of whether or not the neighboring interface is a trunk interface.
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. The default switchport mode for all Ethernet interfaces is dynamic auto .
switchport mode dynamic desirable	Makes the interface actively attempt to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to <i>trunk</i> , <i>desirable</i> , or <i>auto</i> mode.
switchport mode trunk	Puts the interface into permanent trunking mode and negotiates to convert the neighboring 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.
switchport mode dot1q-tunnel	Configures the interface as a tunnel (nontrunking) port to be connected in an asymmetric link with an 802.1Q trunk port. 802.1Q tunneling is used to maintain customer VLAN integrity across a service provider network. See Chapter 13, “Configuring IEEE 802.1Q and Layer 2 Protocol Tunneling,” for more information on tunnel ports.

Encapsulation Types

[Table 10-5](#) lists the Ethernet trunk encapsulation types and keywords.

Table 10-5 Ethernet Trunk Encapsulation Types

Encapsulation	Function
switchport trunk encapsulation isl	Specifies ISL encapsulation on the trunk link.
switchport trunk encapsulation dot1q	Specifies 802.1Q encapsulation on the trunk link.
switchport trunk encapsulation negotiate	Specifies that the interface negotiate with the neighboring interface to become an ISL (preferred) or 802.1Q trunk, depending on the configuration and capabilities of the neighboring interface. This is the default for the switch.
switchport trunk dot1q ethertype value	Sets the ethertype value for 802.1Q encapsulation. Used to select a nonstandard (nondefault) 2-byte ethertype to identify 802.1Q tagged frames. The default ethertype value is 0x8100. This option is supported only on enhanced-services (ES) ports.



Note

The switch does not support Layer 3 trunks; you cannot configure subinterfaces or use the **encapsulation** keyword on Layer 3 interfaces. The switch does support Layer 2 trunks and Layer 3 VLAN interfaces, which provide equivalent capabilities.

The trunking mode, the trunk encapsulation type, and the hardware capabilities of the two connected interfaces determine whether a link becomes an ISL or 802.1Q trunk.

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.
- On an ES port, you can use the **switchport trunk dot1q ethertype value** interface configuration command to program a custom ethertype value for the interface. This feature allows the switch to interoperate with third-party vendor switches that do not use the standard 0x8100 ethertype to identify 802.1Q-tagged frames. For example, if you set 0x1234 as the custom 802.1Q ethertype on a trunk port, incoming frames containing that ethertype are assigned to the VLAN contained in the tag following the ethertype as they would be with a standard 802.1Q trunk. Frames arriving on a trunk port with either the standard ethertype (0x8100) or the custom ethertype value are treated as valid 802.1Q traffic; frames containing any other ethertype are assigned to the native VLAN of that trunk.

Default Layer 2 Ethernet Interface VLAN Configuration

Table 10-6 shows the default Layer 2 Ethernet interface VLAN configuration.

Table 10-6 Default Layer 2 Ethernet Interface VLAN Configuration

Feature	Default Setting
Interface mode	switchport mode dynamic auto
Trunk encapsulation	switchport trunk encapsulation negotiate
802.1Q encapsulation ethertype value	0x8100
Allowed VLAN range	VLANs 1 to 4094
VLAN range eligible for pruning	VLANs 2 to 1001
Default VLAN (for access ports)	VLAN 1
Native VLAN (for 802.1Q trunks)	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, page 10-20](#)
- [Defining the Allowed VLANs on a Trunk, page 10-22](#)
- [Changing the Pruning-Eligible List, page 10-23](#)
- [Configuring the Native VLAN for Untagged Traffic, page 10-24](#)



Note

By default, an interface is in Layer 2 mode. The default mode for Layer 2 interfaces is **switchport mode dynamic auto**. If the neighboring interface supports trunking and is configured to allow trunking, the link is a Layer 2 trunk or, if the interface is in Layer 3 mode, it becomes a Layer 2 trunk when you enter the **switchport** interface configuration command. By default, trunks negotiate encapsulation. If the neighboring interface supports ISL and 802.1Q encapsulation and both interfaces are set to negotiate the encapsulation type, the trunk uses ISL encapsulation.

Interaction with Other Features

Trunking interacts with other features in these ways:

- A trunk port cannot be a secure port.
- A trunk port cannot be a tunnel 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 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.
- We recommend that you configure no more than 24 trunk ports in PVST mode and no more than 40 trunk ports in MST mode.
- 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.

Configuring a Trunk Port

Beginning in privileged EXEC mode, follow these steps to configure a port as an ISL or 802.1Q trunk port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter the interface configuration mode and the port to be configured for trunking.
Step 3	switchport trunk encapsulation {isl dot1q negotiate}	Configure the port to support ISL or 802.1Q encapsulation or to negotiate (the default) with the neighboring interface for encapsulation type. You must configure each end of the link with the same encapsulation type.
Step 4	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 tunnel port or to specify the trunking mode). dynamic auto—Set the interface to a trunk link if the neighboring interface is set to trunk or desirable mode. This is the default. 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.
Step 5	switchport access vlan <i>vlan-id</i>	(Optional) Specify the default VLAN, which is used if the interface stops trunking.
Step 6	switchport trunk native vlan <i>vlan-id</i>	Specify the native VLAN for 802.1Q trunks.
Step 7	end	Return to privileged EXEC mode.
Step 8	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.
Step 9	show interfaces <i>interface-id</i> trunk	Display the trunk configuration of the interface.
Step 10	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 a port 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 gigabitethernet1/0/2
Switch(config-if)# switchport mode dynamic desirable
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# end
```

Defining the Allowed VLANs on a Trunk

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.



Note

VLAN 1 is the default VLAN on all trunk ports in all Cisco switches, and it has previously been a requirement that VLAN 1 always be enabled on every trunk link. You can use the VLAN 1 minimization feature to disable VLAN 1 on any individual VLAN trunk link so that no user traffic (including spanning tree advertisements) is sent or received on VLAN 1.

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. 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, the port will be 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 ISL or 802.1Q trunk:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode and the port to be configured.
Step 3	switchport mode trunk	Configure the interface as a VLAN trunk port.
Step 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.
Step 5	end	Return to privileged EXEC mode.
Step 6	show interfaces <i>interface-id</i> switchport	Verify your entries in the <i>Trunking VLANs Enabled</i> field of the display.
Step 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 on a port:

```
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# switchport trunk allowed vlan remove 2
Switch(config-if)# end
```

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 11-14](#) 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and select the trunk port for which VLANs should be pruned.
Step 3	switchport trunk pruning vlan { add except none remove } <i>vlan-list</i> [<i>,vlan[,vlan[,...]]</i>]	Configure the list of VLANs allowed to be pruned from the trunk. (See the “VTP Pruning” section on page 11-4). 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. 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces <i>interface-id</i> switchport	Verify your entries in the <i>Pruning VLANs Enabled</i> field of the display.
Step 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.

For information about 802.1Q configuration issues, see the [“802.1Q Configuration Considerations” section on page 10-19](#).

Beginning in privileged EXEC mode, follow these steps to configure the native VLAN on an 802.1Q trunk:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and define the interface that is configured as the 802.1Q trunk.
Step 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces <i>interface-id</i> switchport	Verify your entries in the <i>Trunking Native Mode VLAN</i> field.
Step 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.

Configuring Trunk Ports for Load Sharing

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 14, “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 10-3 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 16 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 16 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 10-3 Load Sharing by Using STP Port Priorities



Beginning in privileged EXEC mode, follow these steps to configure the network shown in Figure 10-3.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode on Switch A.
Step 2	vtp domain <i>domain-name</i>	Configure a VTP administrative domain. The domain name can be from 1 to 32 characters.
Step 3	vtp mode server	Configure Switch A as the VTP server.
Step 4	end	Return to privileged EXEC mode.
Step 5	show vtp status	Verify the VTP configuration on both Switch A and Switch B. In the display, check the <i>VTP Operating Mode</i> and the <i>VTP Domain Name</i> fields.
Step 6	show vlan	Verify that the VLANs exist in the database on Switch A.
Step 7	configure terminal	Enter global configuration mode.


	Command	Purpose
Step 8	interface gigabitethernet1/ 0/1	Enter interface configuration mode, and define Gigabit Ethernet port 1 as the interface to be configured as a trunk.
Step 9	switchport trunk encapsulation {isl dot1q negotiate}	Configure the port to support ISL or 802.1Q encapsulation or to negotiate with the neighboring interface. You must configure each end of the link with the same encapsulation type.
Step 10	switchport mode trunk	Configure the port as a trunk port.
Step 11	end	Return to privileged EXEC mode.
Step 12	show interfaces gigabitethernet1/ 0/1 switchport	Verify the VLAN configuration.
Step 13		Repeat Steps 7 through 11 on Switch A for a second interface in the switch.
Step 14		Repeat Steps 7 through 11 on Switch B to configure the trunk ports that connect to the trunk ports configured on Switch A.
Step 15	show vlan	When the trunk links come up, VTP passes the VTP and VLAN information to Switch B. Verify that Switch B has learned the VLAN configuration.
Step 16	configure terminal	Enter global configuration mode on Switch A.
Step 17	interface gigabitethernet1/ 0/1	Enter interface configuration mode, and define the interface to set the STP port priority.
Step 18	spanning-tree vlan 8-10 port-priority 16	Assign the port priority of 16 for VLANs 8 through 10.
Step 19	exit	Return to global configuration mode.
Step 20	interface gigabitethernet1/0/2	Enter interface configuration mode, and define the interface to set the STP port priority.
Step 21	spanning-tree vlan 3-6 port-priority 16	Assign the port priority of 16 for VLANs 3 through 6.
Step 22	end	Return to privileged EXEC mode.
Step 23	show running-config	Verify your entries.
Step 24	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, blocking different ports for different VLANs. The VLANs keep the traffic separate and maintain redundancy in the event of a lost link.

In [Figure 10-4](#), Trunk ports 1 and 2 are configured as 100BASE-T ports. These VLAN path costs are assigned:

- VLANs 2 through 4 are assigned a path cost of 30 on Trunk port 1.
- VLANs 8 through 10 retain the default 100BASE-T path cost on Trunk port 1 of 19.
- VLANs 8 through 10 are assigned a path cost of 30 on Trunk port 2.
- VLANs 2 through 4 retain the default 100BASE-T path cost on Trunk port 2 of 19.

Figure 10-4 Load-Sharing Trunks with Traffic Distributed by Path Cost

Beginning in privileged EXEC mode, follow these steps to configure the network shown in [Figure 10-4](#):

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode on Switch A.
Step 2	interface gigabitethernet1/0/1	Enter interface configuration mode, and define Gigabit Ethernet port 1 as the interface to be configured as a trunk.
Step 3	switchport trunk encapsulation {isl dot1q negotiate}	Configure the port to support ISL or 802.1Q encapsulation. You must configure each end of the link with the same encapsulation type.
Step 4	switchport mode trunk	Configure the port as a trunk port. The trunk defaults to ISL trunking.
Step 5	exit	Return to global configuration mode.
Step 6		Repeat Steps 2 through 4 on a second interface in Switch A.
Step 7	end	Return to privileged EXEC mode.
Step 8	show running-config	Verify your entries. In the display, make sure that the interfaces configured in Steps 2 and 6 are configured as trunk ports.
Step 9	show vlan	When the trunk links come up, Switch A receives the VTP information from the other switches. Verify that Switch A has learned the VLAN configuration.
Step 10	configure terminal	Enter global configuration mode.
Step 11	interface gigabitethernet1/0/1	Enter interface configuration mode, and define Gigabit Ethernet port 1 as the interface on which to set the STP cost.
Step 12	spanning-tree vlan 2-4 cost 30	Set the spanning-tree path cost to 30 for VLANs 2 through 4.
Step 13	end	Return to global configuration mode.
Step 14		Repeat Steps 9 through 11 on the other configured trunk interface on Switch A, and set the spanning-tree path cost to 30 for VLANs 8, 9, and 10.
Step 15	exit	Return to privileged EXEC mode.
Step 16	show running-config	Verify your entries. In the display, verify that the path costs are set correctly for both trunk interfaces.
Step 17	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Configuring VMPS

The VLAN Query Protocol (VQP) is used to support dynamic-access ports, which are not permanently assigned to a VLAN, but given VLAN assignments based on the MAC source addresses seen on the port. Each time an unknown MAC address is seen, the switch sends a VQP query to a remote VMPS; the query includes the newly seen MAC address and the port on which it was seen. The VMPS responds with a VLAN assignment for the port. The switch cannot be a VMPS server but can act as a client to the VMPS and communicate with it through VQP.

This section includes this information about configuring VMPS:

- [“Understanding VMPS” section on page 10-28](#)
- [“Default VMPS Client Configuration” section on page 10-29](#)
- [“VMPS Configuration Guidelines” section on page 10-29](#)
- [“Configuring the VMPS Client” section on page 10-30](#)
- [“Monitoring the VMPS” section on page 10-33](#)
- [“Troubleshooting Dynamic-Access Port VLAN Membership” section on page 10-33](#)
- [“VMPS Configuration Example” section on page 10-34](#)

Understanding VMPS

Each time the client switch receives the MAC address of a new host, it sends a VQP query to the VMPS. When the VMPS receives this query, 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 open or secure mode. In secure mode, the server shuts down the port when an illegal host is detected. In open mode, the server simply denies the host access to the port.

If the port is currently *unassigned* (that is, it does not yet have a VLAN assignment), the VMPS provides one of these responses:

- If the host is allowed on the port, the VMPS sends the client a *vlan-assignment* response containing the assigned VLAN name and allowing access to the host.
- If the host is not allowed on the port and the VMPS is in open 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 port already has a VLAN assignment, the VMPS provides one of these responses:

- If the VLAN in the database matches the current VLAN on the port, the VMPS sends an *success* response, allowing access to the host.
- 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 to and from the host MAC address. The switch continues to monitor the packets directed to the port and sends a query to the VMPS when it identifies a new host 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 or SNMP.

Dynamic-Access Port VLAN Membership

A dynamic-access port can belong to only one VLAN with an ID from 1 to 4094. 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-access 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-access port if they are all in the same VLAN; however, the VMPS shuts down a dynamic-access port if more than 20 hosts are active on the port.

If the link goes down on a dynamic-access 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.

Dynamic-access ports can be used for direct host connections, or they can connect to a network. A maximum of 20 MAC addresses are allowed per port on the switch. A dynamic-access port can belong to only one VLAN at a time, but the VLAN can change over time, depending on the MAC addresses seen.

Default VMPS Client Configuration

Table 10-7 shows the default VMPS and dynamic-access port configuration on client switches.

Table 10-7 Default VMPS Client and Dynamic-Access Port Configuration

Feature	Default Setting
VMPS domain server	None
VMPS reconfirm interval	60 minutes
VMPS server retry count	3
Dynamic-access 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-access ports.
- 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 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.
- A dynamic-access port can participate in fallback bridging.
- The VTP management domain of the VMPS client and the VMPS server must be the same.
- 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vmps server ipaddress primary	Enter the IP address of the switch acting as the primary VMPS server.
Step 3	vmps server ipaddress	(Optional) Enter the IP address of the switch acting as a secondary VMPS server. You can enter up to three secondary server addresses.
Step 4	end	Return to privileged EXEC mode.
Step 5	show vmps	Verify your entries in the <i>VMPS Domain Server</i> field of the display.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.



Note

You must have IP connectivity to the VMPS for dynamic-access ports to work. You can test for IP connectivity by pinging the IP address of the VMPS and verifying that you get a response.

Configuring Dynamic-Access Ports on VMPS Clients

If you are configuring a port on a cluster member switch as a dynamic-access port, first use the **rcommand** privileged EXEC command to log into the cluster member switch.



Caution

Dynamic-access 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode for the switch port that is connected to the end station.
Step 3	switchport mode access	Set the port to access mode.
Step 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.
Step 5	end	Return to privileged EXEC mode.
Step 6	show interfaces <i>interface-id</i> switchport	Verify your entries in the <i>Operational Mode</i> field of the display.
Step 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 auto), 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 vlan** interface configuration command.

Reconfirming VLAN Memberships

Beginning in privileged EXEC mode, follow these steps to confirm the dynamic-access port VLAN membership assignments that the switch has received from the VMPS:

	Command	Purpose
Step 1	vmmps reconfirm	Reconfirm dynamic-access port VLAN membership.
Step 2	show vmmps	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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vmps reconfirm <i>minutes</i>	Enter the number of minutes between reconfirmations of the dynamic VLAN membership. The range is from 1 to 120. The default is 60 minutes.
Step 3	end	Return to privileged EXEC mode.
Step 4	show vmps	Verify the dynamic VLAN reconfirmation status in the <i>Reconfirm Interval</i> field of the display.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vmps retry <i>count</i>	Change the retry count. The retry range is from 1 to 10; the default is 3.
Step 3	end	Return to privileged EXEC mode.
Step 4	show vmps	Verify your entry in the <i>Server Retry Count</i> field of the display.
Step 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 *current*. The one marked *primary* 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 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:          other
```

Troubleshooting Dynamic-Access Port VLAN Membership

The VMPS shuts down a dynamic-access 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-access port.


To re-enable a disabled dynamic-access port, enter the **shutdown** interface configuration command followed by the **no shutdown** interface configuration command.

VMPS Configuration Example

Figure 10-5 shows a network with a VMPS server switch and VMPS client switches with dynamic-access ports. In this example, these assumptions apply:

- The VMPS server and the VMPS client are separate switches.
- The Catalyst 6500 series Switch A is the primary VMPS server.
- The Catalyst 6500 series Switch C and Switch J are secondary VMPS servers.
- End stations are connected to the clients, Switch B and Switch I.
- The database configuration file is stored on the TFTP server with the IP address 172.20.22.7.

Figure 10-5 Dynamic Port VLAN Membership Configuration





Configuring VTP

This chapter describes how to use the VLAN Trunking Protocol (VTP) and the VLAN database for managing VLANs with the Catalyst 3750 Metro 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, page 11-1](#)
- [Configuring VTP, page 11-6](#)
- [Monitoring VTP, page 11-16](#)

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 is designed to work in an environment where updates are made on a single switch and are sent through VTP to other switches in the domain. It does not work well in a situation where multiple updates to the VLAN database occur simultaneously on switches in the same domain, which would result in an inconsistency in the VLAN database.

The switch supports 1005 VLANs, but the number of routed ports, SVIs, and other configured features affects the usage of the switch hardware. If the switch is notified by VTP of a new VLAN and the switch is already using the maximum available hardware resources, it sends a message that there are not enough hardware resources available and shuts down the VLAN. The output of the **show vlan** user EXEC command shows the VLAN in a suspended state.

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 and characteristics.

- [The VTP Domain, page 11-2](#)
- [VTP Modes, page 11-3](#)
- [VTP Advertisements, page 11-3](#)
- [VTP Version 2, page 11-4](#)
- [VTP Pruning, page 11-4](#)

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) 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 11-15](#) 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 trunk connections, including Inter-Switch Link (ISL) and IEEE 802.1Q. VTP dynamically maps VLANs with unique names and internal index associates across multiple LAN types. 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 11-8](#).

VTP Modes

You can configure a supported switch to be in one of the VTP modes listed in [Table 11-1](#).

Table 11-1 VTP Modes

VTP Mode	Description
VTP server	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. In VTP server mode, VLAN configurations are saved in NVRAM. VTP server is the default mode.
VTP client	A VTP client behaves like a VTP server and transmits and receives VTP updates on its trunks, but you cannot create, change, or delete VLANs on a VTP client. VLANs are configured on another switch in the domain that is in server mode. In VTP client mode, VLAN configurations are not saved in NVRAM.
VTP transparent	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 10-12. 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.

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 another switch. Otherwise, the switch cannot receive any VTP advertisements. For more information on trunk ports, see the [“Configuring VLAN Trunks”](#) section on page 10-16.

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 (ISL and 802.1Q)
- 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 10-4](#).
- 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 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 11-1](#) shows a switched network without VTP pruning enabled. Port 1 on Switch A and Port 2 on Switch D are assigned to the Red VLAN. If a broadcast is sent from the host connected to Switch A, Switch A floods the broadcast and every switch in the network receives it, even though Switches C, E, and F have no ports in the Red VLAN.



Figure 11-1 Flooding Traffic without VTP Pruning

Figure 11-2 shows a switched network with VTP pruning enabled. The broadcast traffic from Switch A is not forwarded to Switches C, E, and F because traffic for the Red VLAN has been pruned on the links shown (Port 5 on Switch B and Port 4 on Switch D).

Figure 11-2 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 trunk only (not on all switches in the VTP domain).

See the [“Enabling VTP Pruning” section on page 11-14](#). 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 10-23). 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, page 11-6](#)
- [VTP Configuration Options, page 11-7](#)
- [VTP Configuration Guidelines, page 11-8](#)
- [Configuring a VTP Server, page 11-9](#)
- [Configuring a VTP Client, page 11-11](#)
- [Disabling VTP \(VTP Transparent Mode\), page 11-12](#)
- [Enabling VTP Version 2, page 11-13](#)
- [Enabling VTP Pruning, page 11-14](#)
- [Adding a VTP Client Switch to a VTP Domain, page 11-15](#)

Default VTP Configuration

[Table 11-2](#) shows the default VTP configuration.

Table 11-2 Default VTP Configuration

Feature	Default Setting
VTP domain name	Null.
VTP mode	Server.
VTP version	Version 1 (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, page 11-7](#)
- [VTP Configuration in VLAN Database Configuration Mode, page 11-7](#)

You access VLAN database 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 Database Configuration Mode

You can configure all VTP parameters in VLAN database configuration mode, which you access by entering the **vlan database** privileged EXEC command. For more information about available keywords, refer to the **vtp** VLAN database configuration command description in the command reference for this release. When you enter the **exit** command in VLAN database 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 to and from other switches in the domain.

For more information, see the [“Configuring VLAN Trunks” section on page 10-16](#).

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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vtp mode server	Configure the switch for VTP server mode (the default).
Step 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.
Step 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.
Step 5	end	Return to privileged EXEC mode.
Step 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 database configuration mode to configure VTP parameters.

Beginning in privileged EXEC mode, follow these steps to use VLAN database configuration mode to configure the switch as a VTP server:

	Command	Purpose
Step 1	vlan database	Enter VLAN database configuration mode.
Step 2	vtp server	Configure the switch for VTP server mode (the default).
Step 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.
Step 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 5	exit	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 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 database configuration command.

This example shows how to use VLAN database 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...
Switch#
```

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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vtp mode client	Configure the switch for VTP client mode. The default setting is VTP server.
Step 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.
Step 4	vtp password <i>password</i>	(Optional) Enter the password for the VTP domain.
Step 5	end	Return to privileged EXEC mode.
Step 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** privileged EXEC 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 database configuration mode and entering the **vtp client** command, similar to the second procedure under “[Configuring a VTP Server](#)” section on page 11-9. Use the **no vtp client** VLAN database configuration command to return the switch to VTP server mode or the **no vtp password** VLAN database 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 (the default).

Beginning in privileged EXEC mode, follow these steps to configure VTP transparent mode and save the VTP configuration in the switch startup configuration file:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vtp mode transparent	Configure the switch for VTP transparent mode (disable VTP).
Step 3	end	Return to privileged EXEC mode.
Step 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.
Step 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 the 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 database configuration mode and by entering the **vtp transparent** command, similar to the second procedure under the “[Configuring a VTP Server](#)” section on page 11-9. Use the **no vtp transparent** VLAN database 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 11-8](#).

Beginning in privileged EXEC mode, follow these steps to enable VTP version 2:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vtp version 2	Enable VTP version 2 on the switch. VTP version 2 is disabled by default on VTP version 2-capable switches.
Step 3	end	Return to privileged EXEC mode.
Step 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 database configuration mode and entering the **vtp v2-mode** VLAN database configuration command. To disable VTP version 2, use the **no vtp v2-mode** VLAN database 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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.
Step 3	end	Return to privileged EXEC mode.
Step 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 database configuration mode and entering the **vtp pruning** VLAN database configuration command. To disable VTP pruning, use the **no vtp pruning** VLAN database configuration command. You can also enable VTP version 2 by using the **vtp pruning** privileged EXEC command. However, this command will not be available in future releases.

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 10-23](#).

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:

	Command	Purpose
Step 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 switch configuration revision number.
Step 2	configure terminal	Enter global configuration mode.
Step 3	vtp domain <i>domain-name</i>	Change the domain name from the original one displayed in Step 1 to a new name.
Step 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.
Step 5	show vtp status	Verify that the configuration revision number has been reset to 0.
Step 6	configure terminal	Enter global configuration mode.
Step 7	vtp domain <i>domain-name</i>	Enter the original domain name on the switch.
Step 8	end	The VLAN information on the switch is updated, and you return to privileged EXEC mode.
Step 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 database 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 database 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 11-3](#) shows the privileged EXEC commands for monitoring VTP activity.

Table 11-3 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.



Configuring Voice VLAN

This chapter describes how to configure the voice VLAN feature on the Catalyst 3750 Metro switch. Voice VLAN is also referred to as an *auxiliary VLAN* in some Catalyst 6500 family switch documentation.



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 Voice VLAN, page 12-1](#)
- [Configuring Voice VLAN, page 12-3](#)
- [Displaying Voice VLAN, page 12-6](#)

Understanding Voice VLAN


The voice VLAN feature enables access ports to carry IP voice traffic from an IP phone. When the switch is connected to a Cisco 7960 IP Phone, the IP Phone sends voice traffic with Layer 3 IP precedence and Layer 2 class of service (CoS) values, which are both set to 5 by default. Because the sound quality of an IP phone call can deteriorate if the data is unevenly sent, the switch supports quality of service (QoS) based on IEEE 802.1p CoS. QoS uses classification and scheduling to send network traffic from the switch in a predictable manner. For more information on QoS, see [Chapter 26, “Configuring QoS.”](#)

The Cisco 7960 IP Phone is a configurable device, and you can configure it to forward traffic with an 802.1p priority. You can configure the switch to trust or override the traffic priority assigned by an IP Phone.

The Cisco IP Phone contains an integrated three-port 10/100 switch as shown in [Figure 12-1](#). The ports provide dedicated connections to these devices:

- Port 1 connects to the switch or other voice-over-IP (VoIP) device.
- Port 2 is an internal 10/100 interface that carries the IP phone traffic.
- Port 3 (access port) connects to a PC or other device.

[Figure 12-1](#) shows one way to connect a Cisco 7960 IP Phone.

Figure 12-1 Cisco 7960 IP Phone Connected to a Switch

Cisco IP Phone Voice Traffic

You can configure an access port with an attached Cisco IP Phone to use one VLAN for voice traffic and another VLAN for data traffic from a device attached to the phone. You can configure access ports on the switch to send Cisco Discovery Protocol (CDP) packets that instruct an attached Cisco IP Phone to send voice traffic to the switch in any of these ways:

- In the voice VLAN tagged with a Layer 2 CoS priority value
- In the access VLAN tagged with a Layer 2 CoS priority value
- In the access VLAN, untagged (no Layer 2 CoS priority value)



Note

In all configurations, the voice traffic carries a Layer 3 IP precedence value (the default is 5 for voice traffic and 3 for voice control traffic).

Cisco IP Phone Data Traffic

The switch can also process tagged data traffic (traffic in 802.1Q or 802.1p frame types) from the device attached to the access port on the Cisco IP Phone (see [Figure 12-1](#)). You can configure Layer 2 access ports on the switch to send CDP packets that instruct the attached Cisco IP Phone to configure the IP phone access port in one of these modes:

- In trusted mode, all traffic received through the access port on the Cisco IP Phone passes through the IP phone unchanged.
- In untrusted mode, all traffic in 802.1Q or 802.1p frames received through the access port on the IP phone receive a configured Layer 2 CoS value. The default Layer 2 CoS value is 0. Untrusted mode is the default.



Note

Untagged traffic from the device attached to the Cisco IP Phone passes through the IP phone unchanged, regardless of the trust state of the access port on the IP phone.

Configuring Voice VLAN

This section describes how to configure voice VLAN on access ports. This section contains this configuration information:

- [Default Voice VLAN Configuration, page 12-3](#)
- [Voice VLAN Configuration Guidelines, page 12-3](#)
- [Configuring a Port Connected to a Cisco 7960 IP Phone, page 12-4](#)

Default Voice VLAN Configuration

The voice VLAN feature is disabled by default.

When the voice VLAN feature is enabled, all untagged traffic is sent according to the default CoS priority of the port.

The CoS value is not trusted for 802.1p or 802.1Q tagged traffic.

Voice VLAN Configuration Guidelines

These are the voice VLAN configuration guidelines:

- You should configure voice VLAN on switch access ports; voice VLAN is not supported on trunk ports. You can only configure a voice VLAN on Layer 2 ports.

**Note**

Voice VLAN is only supported on access ports and not on trunk ports, even though the configuration is allowed.

- Before you enable voice VLAN, we recommend that you enable QoS on the switch by entering the **mls qos** global configuration command and configure the port trust state to trust by entering the **mls qos trust cos** interface configuration command. If you use the auto-QoS feature, these settings are automatically configured. For more information, see [Chapter 26, “Configuring QoS.”](#)
- You must enable CDP on the switch port connected to the Cisco IP Phone to send configuration to the Cisco IP Phone. (CDP is enabled by default globally and on all switch interfaces.)
- The Port Fast feature is automatically enabled when voice VLAN is configured. When you disable voice VLAN, the Port Fast feature is not automatically disabled.
- If the Cisco IP Phone and a device attached to the Cisco IP Phone are in the same VLAN, they must be in the same IP subnet. These conditions indicate that they are in the same VLAN:
 - They both use 802.1p or untagged frames.
 - The Cisco IP Phone uses 802.1p frames and the device uses untagged frames.
 - The Cisco IP Phone uses untagged frames and the device uses 802.1p frames.
 - The Cisco IP Phone uses 802.1Q frames and the voice VLAN is the same as the access VLAN.
- The Cisco IP Phone and a device attached to the phone cannot communicate if they are in the same VLAN and subnet but use different frame types because traffic in the same subnet is not routed (routing would eliminate the frame type difference).
- You cannot configure static secure MAC addresses in the voice VLAN.

- Voice VLAN ports can also be these port types:
 - Dynamic access port. See the “[Configuring Dynamic-Access Ports on VMPS Clients](#)” section on page 10-31 for more information.
 - 802.1X authenticated port. See the “[Configuring 802.1x Authentication](#)” section on page 8-11 for more information.
 - Protected port. See the “[Configuring Protected Ports](#)” section on page 18-4 for more information.
 - A source or destination port for a SPAN or RSPAN session.
 - Secure port. See the “[Configuring Port Security](#)” section on page 18-6 for more information.

**Note**

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 two plus the maximum number of secure addresses allowed on the access VLAN. When the port is connected to a Cisco IP phone, the IP phone requires up to two MAC addresses. The IP phone address is learned on the voice VLAN and might also be learned on the access VLAN. Connecting a PC to the IP phone requires additional MAC addresses.

Configuring a Port Connected to a Cisco 7960 IP Phone

Because a Cisco 7960 IP Phone also supports a connection to a PC or other device, a port connecting the switch to a Cisco IP Phone can carry mixed traffic. You can configure a port to determine how the IP phone carries voice traffic and data traffic.

This section includes these topics:

- [Configuring IP Phone Voice Traffic, page 12-4](#)
- [Configuring the Priority of Incoming Data Frames, page 12-5](#)

Configuring IP Phone Voice Traffic

You can configure a port connected to the Cisco IP Phone to send CDP packets to the phone to configure the way in which the phone sends voice traffic. The phone can carry voice traffic in 802.1Q frames for a specified voice VLAN with a Layer 2 CoS value. It can use 802.1P priority tagging to give voice traffic a higher priority and forward all voice traffic through the native (access) VLAN. The IP phone can also send untagged voice traffic or use its own configuration to send voice traffic in the access VLAN. In all configurations, the voice traffic carries a Layer 3 IP precedence value (the default is 5).

Beginning in privileged EXEC mode, follow these steps to configure voice traffic on a port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface connected to the IP phone.
Step 3	mls qos trust cos	Configure the interface to classify ingress traffic packets by using the packet CoS value. For untagged packets, the port default CoS value is used.
		Note Before configuring the port trust state, you must first globally enable QoS by using the mls qos global configuration command.

	Command	Purpose
Step 4	switchport voice vlan { <i>vlan-id</i> dot1p none untagged }	Configure how the Cisco IP Phone carries voice traffic: vlan-id—Configure the Cisco IP Phone to forward all voice traffic through the specified VLAN. By default, the Cisco IP Phone forwards the voice traffic with an 802.1Q priority of 5. Valid VLAN IDs are from 1 to 4094. dot1p—Configure the Cisco IP Phone to use 802.1P priority tagging for voice traffic and to use the default native VLAN (VLAN 0) to carry all traffic. By default, the Cisco IP Phone forwards the voice traffic with an 802.1P priority of 5. none—Allow the IP phone to use its own configuration to send untagged voice traffic. untagged—Configure the phone to send untagged voice traffic.
Step 5	end	Return to privileged EXEC mode.
Step 6	show interfaces <i>interface-id</i> switchport or show running-config interface <i>interface-id</i>	Verify your voice VLAN entries. Verify your QoS and voice VLAN entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to configure a port connected to an IP phone to use the CoS value to classify ingress traffic, to use 802.1P priority tagging for voice traffic, and to use and the default native VLAN (VLAN 0) to carry all traffic:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# mls qos trust cos
Switch(config-if)# switchport voice vlan dot1p
Switch(config-if)# end
```

To return the port to its default setting, use the **no switchport voice vlan** interface configuration command.

Configuring the Priority of Incoming Data Frames

You can connect a PC or other data device to a Cisco IP Phone port. To process tagged data traffic (in 802.1Q or 802.1P frames), you can configure the switch to send CDP packets to instruct the IP phone how to send data packets from the device attached to the access port on the Cisco IP Phone. The PC can generate packets with an assigned CoS value. You can configure the Cisco IP Phone to not change (trust) or to override (not trust) the priority of frames arriving on the IP phone port from connected devices.

Beginning in privileged EXEC mode, follow these steps to set the priority of data traffic received from the nonvoice port on the Cisco IP Phone:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface connected to the IP phone.

	Command	Purpose
Step 3	switchport priority extend { <i>cos value</i> trust }	Set the priority of data traffic received from the IP phone access port: • cos value—Configure the IP phone to override the priority received from the PC or the attached device with the specified CoS value. The value is a number from 0 to 7, with 7 as the highest priority. The default priority is cos 0. • trust—Configure the IP phone access port to trust the priority received from the PC or the attached device.
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces <i>interface-id</i> switchport	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to configure a port connected to an IP phone to not change the priority of frames received from the PC or the attached device:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# switchport priority extend trust
Switch(config-if)# end
```

To return the port to its default setting, use the **no switchport priority extend** interface configuration command.

Displaying Voice VLAN

To display voice VLAN configuration for an interface, use the **show interfaces** *interface-id* **switchport** privileged EXEC command.



Configuring IEEE 802.1Q and Layer 2 Protocol Tunneling

Virtual private networks (VPNs) provide enterprise-scale connectivity on a shared infrastructure, often Ethernet-based, with the same security, prioritization, reliability, and manageability requirements of private networks. Tunneling is a feature designed for service providers who carry traffic of multiple customers across their networks and are required to maintain the VLAN and Layer 2 protocol configurations of each customer without impacting the traffic of other customers. The Catalyst 3750 Metro switch supports 802.1Q tunneling and Layer 2 protocol tunneling, as well as VLAN mapping (VLAN-ID translation).



Note

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

This chapter contains these sections:

- [Understanding 802.1Q Tunneling, page 13-2](#)
- [Configuring 802.1Q Tunneling, page 13-4](#)
- [Configuring VLAN Mapping, page 13-7](#)
- [Understanding Layer 2 Protocol Tunneling, page 13-10](#)
- [Configuring Layer 2 Protocol Tunneling, page 13-12](#)
- [Monitoring and Maintaining Tunneling and Mapping Status, page 13-16](#)



Note

For information about VPNs used with multiprotocol label switching (MPLS), see [Chapter 30](#), “Configuring MPLS and EoMPLS.”


Understanding 802.1Q Tunneling

Service-provider business customers often have specific requirements for VLAN IDs and the number of VLANs to be supported. The VLAN ranges required by different customers in the same service-provider network might overlap, and traffic of customers through the infrastructure might be mixed. Assigning a unique range of VLAN IDs to each customer would restrict customer configurations and could easily exceed the VLAN limit (4096) of the 802.1Q specification.

Using the 802.1Q tunneling feature, service providers can use a single VLAN to support customers who have multiple VLANs. Customer VLAN IDs are preserved, and traffic from different customers is segregated within the service-provider infrastructure, even when they appear to be on the same VLAN. Using 802.1Q tunneling expands VLAN space by using a VLAN-in-VLAN hierarchy and tagging the tagged packets. A port configured to support 802.1Q tunneling is called a *tunnel port*. When you configure tunneling, you assign a tunnel port to a VLAN that is dedicated to tunneling. Each customer requires a separate service-provider VLAN ID, but that VLAN ID supports all of the customer's VLANs.

Customer traffic tagged in the normal way with appropriate VLAN IDs come from an 802.1Q trunk port on the customer device and into a tunnel port on the service-provider edge switch. The link between the customer device and the edge switch is an asymmetric link because one end is configured as an 802.1Q trunk port and the other end is configured as a tunnel port. You assign the tunnel port interface to an access VLAN ID that is unique to each customer. See [Figure 13-1](#).


Figure 13-1 802.1Q Tunnel Ports in a Service-Provider Network



Packets coming from the customer trunk port into the tunnel port on the service-provider edge switch are normally 802.1Q-tagged with the appropriate VLAN ID. The tagged packets remain intact inside the switch and, when they exit the trunk port into the service-provider network, are encapsulated with another layer of an 802.1Q tag (called the *metro tag*) that contains the VLAN ID that is unique to the customer. The original 802.1Q tag from the customer is preserved in the encapsulated packet. Therefore, packets entering the service-provider infrastructure are double-tagged, with the outer tag containing the customer's access VLAN ID, and the inner VLAN ID being the VLAN of the incoming traffic.

When the double-tagged packet enters another trunk port in a service-provider core switch, the outer tag is stripped as the packet is processed inside the switch. When the packet exits another trunk port on the same core switch, the same metro tag is again added to the packet. Figure 13-2 shows the structure of the double-tagged packet.

Figure 13-2 Normal, 802.1Q, and Double-Tagged Ethernet Packet Formats



When the packet enters the trunk port of the service-provider egress switch, the outer tag is again stripped as the packet is processed internally on the switch. However, the metro tag is not added when it is sent out the tunnel port on the edge switch into the customer network, and the packet is sent as a normal 802.1Q-tagged frame to preserve the original VLAN numbers in the customer network.

In Figure 13-1, Customer A was assigned VLAN 30, and Customer B was assigned VLAN 40. Packets entering the edge-switch tunnel ports with 802.1Q tags are double-tagged when they enter the service-provider network, with the outer tag containing VLAN ID 30 or 40, appropriately, and the inner tag containing the original VLAN number, for example, VLAN 100. Even if both Customers A and B have VLAN 100 in their networks, the traffic remains segregated within the service-provider network because the outer tag is different. With 802.1Q tunneling, each customer controls its own VLAN numbering space, which is independent of the VLAN numbering space used by other customers and the VLAN numbering space used by the service-provider network.

At the outbound tunnel port, the original VLAN numbers on the customer's network are recovered. It is possible to have multiple levels of tunneling and tagging, but the switch supports only one level in this release.

If the traffic coming from a customer network is not tagged (native VLAN frames), these packets are bridged or routed as if they were normal packets. All packets entering the service-provider network through a tunnel port on an edge switch are treated as untagged packets, whether they are untagged or already tagged with 802.1Q headers. The packets are encapsulated with the metro tag VLAN ID (set to the access VLAN of the tunnel port) when they are sent through the service-provider network on an 802.1Q trunk port. The priority field on the metro tag is set to the interface class of service (CoS) priority configured on the tunnel port (the default is zero if none is configured).

The switch supports intelligent 802.1Q tunneling QoS, or the ability to copy the inner CoS value to the outer CoS value. For more information, see the [“Configuring the Trust State on Ports Within the QoS Domain” section on page 26-42](#) and the [“Configuring the CoS Value for an Interface” section on page 26-45](#).

Configuring 802.1Q Tunneling

This section includes this information about configuring 802.1Q tunneling:

- [Default 802.1Q Tunneling Configuration, page 13-4](#)
- [802.1Q Tunneling Configuration Guidelines, page 13-4](#)
- [802.1Q Tunneling and Other Features, page 13-6](#)
- [Configuring an 802.1Q Tunneling Port, page 13-6](#)

Default 802.1Q Tunneling Configuration

By default, 802.1Q tunneling is disabled because the default switchport mode is dynamic auto. Tagging of 802.1Q native VLAN packets on all 802.1Q trunk ports is also disabled.

802.1Q Tunneling Configuration Guidelines

When you configure 802.1Q tunneling, you should always use an asymmetrical link between the customer device and the edge switch, with the customer device port configured as an 802.1Q trunk port and the edge switch port configured as a tunnel port.

Assign tunnel ports only to VLANs that are used for tunneling.

Configuration requirements for native VLANs and maximum transmission units (MTUs) are explained in the next sections.

Native VLANs


When configuring 802.1Q tunneling on an edge switch, you must use 802.1Q trunk ports for sending out packets into the service-provider network. However, packets going through the core of the service-provider network might be carried through 802.1Q trunks, ISL trunks, or nontrunking links. When 802.1Q trunks are used in these core switches, the native VLANs of the 802.1Q trunks must not match any native VLAN of the nontrunking (tunneling) port on the same switch because traffic on the native VLAN would not be tagged on the 802.1Q transmitting trunk port.

See [Figure 13-3](#). VLAN 40 is configured as the native VLAN for the 802.1Q trunk port from Customer X at the ingress edge switch in the service-provider network (Switch B). Switch A of Customer X sends a tagged packet on VLAN 30 to the ingress tunnel port of Switch B in the service-provider network, which belongs to access VLAN 40. Because the access VLAN of the tunnel port (VLAN 40) is the same as the native VLAN of the edge switch trunk port (VLAN 40), the metro tag is not added to tagged packets received from the tunnel port. The packet carries only the VLAN 30 tag through the service-provider network to the trunk port of the egress edge switch (Switch C) and is misdirected through the egress switch tunnel port to Customer Y.

These are some ways to solve this problem:

- Use ISL trunks between core switches in the service-provider network. Although customer interfaces connected to edge switches must be 802.1Q trunks, we recommend using ISL trunks for connecting switches in the core layer.
- Use the **vlan dot1q tag native** global configuration command to configure the edge switch so that all packets going out an 802.1Q trunk, including the native VLAN, are tagged. If the switch is configured to tag native VLAN packets on all 802.1Q trunks, the switch accepts untagged packets, but sends only tagged packets.
- Ensure that the native VLAN ID on the edge switch trunk port is not within the customer VLAN range. For example, if the trunk port carries traffic of VLANs 100 to 200, assign the native VLAN a number outside that range.

Figure 13-3 Potential Problem with 802.1Q Tunneling and Native VLANs



System MTU

The default system MTU for traffic on the Catalyst 3750 Metro switch is 1500 bytes. You can configure Fast Ethernet ports to support frames larger than 1500 bytes by using the **system mtu** global configuration command. You can configure Gigabit Ethernet ports to support frames larger than 1500 bytes by using the **system mtu jumbo** global configuration command. Because the 802.1Q tunneling feature increases the frame size by 4 bytes when the metro tag is added, you must configure all switches in the service-provider network to be able to process maximum frames by increasing the switch system MTU size to at least 1504 bytes. The maximum allowable system MTU for Gigabit Ethernet interfaces is 9000 bytes; the maximum system MTU for Fast Ethernet interfaces is 1546 bytes.

802.1Q Tunneling and Other Features

Although 802.1Q tunneling works well for Layer 2 packet switching, there are incompatibilities with some Layer 2 features and with Layer 3 switching.

- A tunnel port cannot be a routed port.
- IP routing is not supported on a VLAN that includes 802.1Q ports. Packets received from a tunnel port are forwarded based only on Layer 2 information. If routing is enabled on the switch virtual interface (SVI) that includes tunnel ports, untagged IP packets received from the tunnel port are recognized and routed by the switch. This allows the customer to access the Internet through its native VLAN. If this access is not required, you should not configure SVIs on VLANs that include tunnel ports.
- Fallback bridging is not supported on tunnel ports. Because all 802.1Q-tagged packets received from a tunnel port are treated as non-IP packets, if fallback bridging is enabled on VLANs that have tunnel ports configured, IP packets would be improperly bridged across VLANs. Therefore, you must *not* enable fallback bridging on VLANs with tunnel ports.
- Tunnel ports do not support IP access control lists (ACLs).
- Layer 3 quality of service (QoS) ACLs and other QoS features related to Layer 3 information are not supported on tunnel ports. MAC-based QoS ACLs are supported on tunnel ports.
- EtherChannel port groups are compatible with tunnel ports as long as the 802.1Q configuration is consistent within an EtherChannel port group.
- Port Aggregation Protocol (PAgP), Link Aggregation Control Protocol (LACP), and UniDirectional Link Detection (UDLD) are supported on 802.1Q tunnel ports.
- Dynamic Trunking Protocol (DTP) is not compatible with 802.1Q tunneling because you must manually configure asymmetric links with tunnel ports and trunk ports.
- VLAN Trunking Protocol (VTP) does not work between devices that are connected by an asymmetrical link or devices that communicate through a tunnel.
- Loopback detection is supported on 802.1Q tunnel ports.
- When a port is configured as an 802.1Q tunnel port, spanning-tree bridge protocol data unit (BPDU) filtering is automatically enabled on the interface. Cisco Discovery Protocol (CDP) is automatically disabled on the interface.

Configuring an 802.1Q Tunneling Port

Beginning in privileged EXEC mode, follow these steps to configure a port as an 802.1Q tunnel port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode and the interface to be configured as a tunnel port. This should be the edge port in the service-provider network that connects to the customer switch. Valid interfaces include physical interfaces and port-channel logical interfaces (port channels 1 to 12).
Step 3	switchport access vlan <i>vlan-id</i>	Specify the default VLAN, which is used if the interface stops trunking. This VLAN ID is specific to the particular customer.
Step 4	switchport mode dot1q-tunnel	Set the interface as an 802.1Q tunnel port.

	Command	Purpose
Step 5	exit	Return to global configuration mode.
Step 6	vlan dot1q tag native	(Optional) Set the switch to enable tagging of native VLAN packets on all 802.1Q trunk ports. When not set, if a customer VLAN ID is the same as the native VLAN, the trunk port does not apply a metro tag, and packets might be sent to the wrong destination.
Step 7	end	Return to privileged EXEC mode.
Step 8	show dot1q-tunnel	Display the tunnel ports on the switch.
Step 9	show vlan dot1q tag native	Display 802.1Q native VLAN tagging status.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no switchport mode dot1q-tunnel** interface configuration command to return the port to the default state of dynamic auto. Use the **no vlan dot1q tag native** global configuration command to disable tagging of native VLAN packets.

This example shows how to configure a port as a tunnel port, enable tagging of native VLAN packets, and verify the configuration. In this configuration, the VLAN ID for the customer connected to Fast Ethernet port 7 is VLAN 22.

```
Switch(config)# interface fastethernet1/0/7
Switch(config-if)# switchport access vlan 22
% Access VLAN does not exist. Creating vlan 22
Switch(config-if)# switchport mode dot1q-tunnel
Switch(config-if)# exit
Switch(config)# vlan dot1q tag native
Switch(config)# end
Switch# show dot1q-tunnel interface fastethernet1/0/7
Port
-----
Fa1/0/7
-----
Switch# show vlan dot1q tag native
dot1q native vlan tagging is enabled
```

Configuring VLAN Mapping

VLAN mapping, or VLAN ID translation, is a feature that you can configure on the enhanced-services (ES) ports connected to the service-provider network to map the customer VLANs to service-provider VLANs. VLAN mapping acts as a filter on the ES ports without affecting the internal operation of the switch or the customer VLANs. A switch might also have a number of reserved VLANs or have a limited VLAN range. When customers want to use a VLAN number in the reserved range, you can use VLAN mapping to overlap customer VLANs by encapsulating the customer traffic in 802.1Q tunnels.

With VLAN mapping, the customer VLAN ID in the 802.1Q or ISL tag, or the inner and outer tag in an 802.1Q tunneled frame, are mapped (or translated) just before a packet is transmitted and just after a packet is received. The service-provider VLANs are not seen by the switch, so all configuration and statistics are done with the customer side-VLANs.



Note

Do not configure VLAN mapping on an interface configured for MPLS or EoMPLS.

This section includes this configuration information:

- [Default VLAN Mapping Configuration, page 13-8](#)
- [Mapping Customer VLANs to Service-Provider VLANs, page 13-8](#)
- [Mapping Customer 802.1Q Traffic with VLAN IDs, page 13-9](#)


Default VLAN Mapping Configuration

By default, no VLAN mapping is configured.

Mapping Customer VLANs to Service-Provider VLANs

Figure 13-4 shows a topology where a customer uses the same VLANs in multiple sites at different sides of a service-provider network. You map the customer VLAN IDs to service-provider VLAN IDs for packet travel across the service-provider backbone. The customer VLAN IDs are retrieved at the other side of the service-provider backbone for use in the other customer site. Configure the same set of VLAN mappings at an ES on each side of the service-provider network. See the example following the configuration steps.

Figure 13-4 Example of VLAN Mapping



Beginning in privileged EXEC mode, follow these steps on an ES port to map a customer VLAN ID to a service-provider VLAN ID:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode and the ES interface connected to the service-provider network.
Step 3	switchport vlan mapping <i>vlan-id translated-id</i>	Enter the VLAN IDs to be mapped: <i>vlan-id</i>—the customer VLAN ID entering the switch from the customer network. The range is from 1 to 4094. <i>translated-id</i>—the assigned service-provider VLAN ID. The range is from 1 to 4094.
Step 4	end	Return to privileged EXEC mode.
Step 5	show vlan mapping	Verify the configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no switchport vlan mapping** *vlan-id translated-id* command to remove the VLAN translation information.

This example shows how to map VLAN IDs 1, 2, 3, 4, and 5 in the customer network to VLANs 2001, 2002, 2003, and 2004 in the service-provider network as shown in [Figure 13-4](#). You configure these same VLAN mapping commands for an ES port in Switch A and Switch B.

```
Switch(config)# interface gigabitEthernet1/1/1
Switch(config-if)# switchport vlan mapping 1 1001
Switch(config-if)# switchport vlan mapping 2 1002
Switch(config-if)# switchport vlan mapping 3 1003
Switch(config-if)# switchport vlan mapping 4 1004
Switch(config-if)# switchport vlan mapping 5 1005
Switch(config-if)# exit
```

Mapping Customer 802.1Q Traffic with VLAN IDs

You can also use the **switchport vlan mapping** interface configuration command with the **dot1qtunnel** keyword to map 802.1Q traffic with a specified outer VLAN ID and inner VLAN ID to a single VLAN ID in the service-provider network. You can use **drop** keyword to specify that traffic on a specified outer VLAN ID is dropped unless the specified inner VLAN and outer VLAN ID combination is explicitly translated.

Beginning in privileged EXEC mode, follow these steps on an ES port to configure the switch to drop or translate 802.1Q tagged traffic:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode and the ES interface connected to the service-provider network.

	Command	Purpose
Step 3	switchport vlan mapping dot1qtunnel <i>outer-vlan-id {inner-vlan-id translated-vlan-id drop}</i>	Specify the action to be taken if the outer VLAN ID of the original 802.1Q packet matches the specified value: <i>outer-vlan-id</i> —Enter the outer VLAN ID of the original 802.1Q packet. <i>inner-vlan-id translated-vlan-id</i> —Enter the inner VLAN ID of the original 802.1Q packet and the translated VLAN ID for the service-provider network. The VLAN ID range is from 1 to 4094. drop—Specify that traffic with the specified outer VLAN ID is dropped unless the inner VLAN ID is matched in the existing VLAN mappings.
Step 4	end	Return to privileged EXEC mode.
Step 5	show vlan mapping	Verify the configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no switchport vlan mapping dot1qtunnel** *outer-vlan-id {inner-vlan-id translated-vlan-id | drop}* interface configuration command to remove the mapping information.

This example shows how to configure the ES port so that 802.1Q traffic with the outer VLAN ID 5 and the inner VLAN ID 6 would leave the switch with the single VLAN ID of 10. The 802.1Q traffic with the outer VLAN 5 and any inner VLAN ID other than 6 would be dropped. Other traffic on VLAN 5 would also be dropped:

```
Switch(config)# interface gigabiethernet1/1/1
Switch(config-if)# switchport vlan mapping dot1q-tunnel 5 6 10
Switch(config-if)# switchport vlan mapping dot1q-tunnel 5 drop
Switch(config-if)# exit
```

Understanding Layer 2 Protocol Tunneling

Customers at different sites connected across a service-provider network need to run various Layer 2 protocols to scale their topology to include all remote sites, as well as the local sites. STP must run properly, and every VLAN should build a proper spanning tree that includes the local site and all remote sites across the service-provider infrastructure. Cisco Discovery Protocol (CDP) must discover neighboring Cisco devices from local and remote sites. VLAN Trunking Protocol (VTP) must provide consistent VLAN configuration throughout all sites in the customer network.

When protocol tunneling is enabled, edge switches on the inbound side of the service-provider infrastructure encapsulate Layer 2 protocol packets with a special MAC address and send them across the service-provider network. Core switches in the network do not process these packets but forward them as normal packets. Layer 2 protocol data units (PDUs) for CDP, STP, or VTP cross the service-provider infrastructure and are delivered to customer switches on the outbound side of the service-provider network. Identical packets are received by all customer ports on the same VLANs with these results:

- Users on each of a customer's sites are able to properly run STP, and every VLAN can build a correct spanning tree based on parameters from all sites and not just from the local site.
- CDP discovers and shows information about the other Cisco devices connected through the service-provider network.

- VTP provides consistent VLAN configuration throughout the customer network, propagating through the service provider to all switches.

**Note**

To provide interoperability with third-party vendors, you can use the Layer 2 protocol-tunnel bypass feature. Bypass mode transparently forwards control PDUs to vendor switches that have different ways of controlling protocol tunneling. You implement bypass mode by enabling Layer 2 protocol tunneling on the egress trunk port.

Layer 2 protocol tunneling can be used independently or can enhance 802.1Q tunneling. If protocol tunneling is not enabled on 802.1Q tunneling ports, remote switches at the receiving end of the service-provider network do not receive the PDUs and cannot properly run STP, CDP, and VTP. When protocol tunneling is enabled, Layer 2 protocols within each customer's network are totally separate from those running within the service-provider network. Customer switches on different sites that send traffic through the service-provider network with 802.1Q tunneling achieve complete knowledge of the customer's VLAN. If 802.1Q tunneling is not used, you can still enable Layer 2 protocol tunneling by connecting to the customer switch through access ports or trunk ports and enabling tunneling on the service-provider access or trunk port.

For example, in [Figure 13-5](#), Customer X has four switches in the same VLAN that are connected through the service-provider network. If the network does not tunnel PDUs, switches on the far ends of the network cannot properly run STP, CDP, and VTP. For example, STP for a VLAN on a switch in Customer X, Site 1 will build a spanning tree on the switches at that site without considering convergence parameters based on Customer X's switch in Site 2. This could result in the topology shown in [Figure 13-6](#).

Figure 13-5 Layer 2 Protocol Tunneling





Figure 13-6 Virtual Network Topology without BPDU Tunneling

Configuring Layer 2 Protocol Tunneling

You enable Layer 2 protocol tunneling (by protocol) on the ports that are connected to the customer in the edge switches of the service-provider network. Edge-switch tunnel ports are connected to customer 802.1Q trunk ports; edge-switch access ports are connected to customer access ports. The edge switches connected to the customer switch perform the tunneling process.

You can enable Layer 2 protocol tunneling on ports that are configured as access ports, tunnel ports, or trunk ports. You cannot enable Layer 2 protocol tunneling on ports configured with switchport mode **dynamic auto** or **dynamic desirable**. The switch supports Layer 2 protocol tunneling for CDP, STP, and VTP.

When the Layer 2 PDUs that entered the inbound edge switch through a Layer 2 protocol-enabled port exit the switch through the trunk port into the service-provider network, the switch overwrites the customer PDU-destination MAC address with a well-known Cisco proprietary multicast address (01-00-0c-cd-cd-d0). If 802.1Q tunneling is enabled, packets are also double-tagged; the outer tag is the customer metro tag and the inner tag is the customer VLAN tag. The core switches ignore the inner tags and forward the packet to all trunk ports in the same metro VLAN. The edge switches on the outbound side restore the proper Layer 2 protocol and MAC address information and forward the packets to all tunnel ports, access ports, and Layer 2 protocol-enabled trunk ports in the same metro VLAN. Therefore, the Layer 2 PDUs are kept intact and delivered across the service-provider infrastructure to the other side of the customer network.

See [Figure 13-5](#), with Customer X and Customer Y in access VLANs 30 and 40, respectively.

Asymmetric links connect the customers in Site 1 to edge switches in the service-provider network. The Layer 2 PDUs (for example, BPDUs) coming into Switch B from Customer Y in Site 1 are forwarded to the infrastructure as double-tagged packets with the well-known MAC address as the destination MAC address. These double-tagged packets have the outer VLAN tag of 40 as well as an inner VLAN tag (for example, VLAN 100). When the double-tagged packets reach Switch D, the outer VLAN tag 40 is removed, the well-known MAC address is replaced with the respective Layer 2 protocol MAC address, and the packet is sent to Customer Y on Site 2 as a single-tagged frame in VLAN 100.

You can also enable Layer 2 protocol tunneling on access or trunk ports on the edge switch connected to access or trunk ports on the customer switch. In this case, the encapsulation and de-encapsulation behavior is the same as described in the previous paragraph, except that the packets are not double-tagged in the service-provider network. The single tag is the customer-specific access VLAN tag.

This section contains this information about configuring Layer 2 protocol tunneling:

- [Default Layer 2 Protocol Tunneling Configuration, page 13-13](#)
- [Layer 2 Protocol Tunneling Configuration Guidelines, page 13-13](#)
- [Configuring Layer 2 Tunneling, page 13-14](#)

Default Layer 2 Protocol Tunneling Configuration

[Table 13-1](#) shows the default Layer 2 protocol tunneling configuration.

Table 13-1 Default Layer 2 Ethernet Interface VLAN Configuration

Feature	Default Setting
Layer 2 protocol tunneling	Disabled for CDP, STP, and VTP.
Shutdown threshold	No threshold for packets-per-second of Layer 2 PDUs per port for the port to shut down.
Drop threshold	No threshold for packets-per-second of Layer 2 PDUs per port for the port to drop the PDUs.
Class of service (CoS) value	If a CoS value is configured on the interface for data packets, that value is the default used for Layer 2 PDUs. If none is configured, the default is 5.

Layer 2 Protocol Tunneling Configuration Guidelines

These are some configuration guidelines and operating characteristics of Layer 2 protocol tunneling:

- The switch supports tunneling of CDP, STP, including multiple STP (MSTP), and VTP. Protocol tunneling is disabled by default but can be enabled for the individual protocols on 802.1Q tunnel ports, access ports, or trunk ports.
- The switch does not support Layer 2 protocol tunneling on ports with switchport mode **dynamic auto** or **dynamic desirable**.
- Dynamic Trunking Protocol (DTP) is not compatible with Layer 2 protocol tunneling because you must manually configure asymmetric links with tunnel ports and trunk ports.
- The edge switches on the outbound side restore the proper Layer 2 protocol and MAC address information and forward the packets to all tunnel ports, access ports and Layer 2 protocol-enabled trunk ports in the same metro VLAN.
- For interoperability with third-party vendor switches, the switch supports a Layer 2 protocol-tunnel bypass feature. Bypass mode transparently forwards control PDUs to vendor switches that have different ways of controlling protocol tunneling. When Layer 2 protocol tunneling is enabled on ingress ports on a switch, egress trunk ports forward the tunneled packets with a special encapsulation. If you also enable Layer 2 protocol tunneling on the egress trunk port, this behavior is bypassed and the switch forwards control PDUs without any processing or modification.
- EtherChannel port groups are compatible with tunnel ports as long as the 802.1Q configuration is consistent within an EtherChannel port group.

- If an encapsulated PDU (with the proprietary destination MAC address) is received from a tunnel port or access or trunk port with Layer 2 tunneling enabled, the port is shut down to prevent loops. The port also shuts down when a configured shutdown threshold for the protocol is reached. You can manually re-enable the port (by entering a **shutdown** and **no shutdown** command sequence) or if errdisable recovery is enabled, the operation is retried after a specified time interval.
- Only decapsulated PDUs are forwarded to the customer network. The spanning-tree instance running on the service-provider network does not forward BPDUs to tunnel ports. No CDP packets are forwarded from tunnel ports.
- When protocol tunneling is enabled on an interface, you can set a per-protocol, per-port, shutdown threshold for the PDUs generated by the customer network. If the limit is exceeded, the port shuts down. You can also rate-limit BPDUs by using QoS ACLs and policy maps on a tunnel port.
- When protocol tunneling is enabled on an interface, you can set a per-protocol, per-port, drop threshold for the PDUs generated by the customer network. If the limit is exceeded, the port drops PDUs until the rate at which the port receives them is below the drop threshold.
- Because tunneled PDUs (especially STP BPDUs) must be delivered to all remote sites for the customer virtual network to operate properly, you can give PDUs higher priority within the service-provider network than data packets received from the same tunnel port. By default, the PDUs use the same CoS value as data packets.

Configuring Layer 2 Tunneling

Beginning in privileged EXEC mode, follow these steps to configure a port for Layer 2 protocol tunneling:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter the interface configuration mode and the interface to be configured as a tunnel port. This should be the edge port in the service-provider network that connects to the customer switch. Valid interfaces include physical interfaces and port-channel logical interfaces (port channels 1 to 12).
Step 3	switchport mode access or switchport mode dot1q-tunnel or switchport mode trunk	Configure the interface as an access port, an 802.1Q tunnel port or a trunk port.
Step 4	l2protocol-tunnel [cdp stp vtp]	Enable protocol tunneling for the desired protocol. If no keyword is entered, tunneling is enabled for all three Layer 2 protocols.
Step 5	l2protocol-tunnel shutdown-threshold [cdp stp vtp] <i>value</i>	(Optional) Configure the threshold in packets per second to be received for encapsulation before the interface shuts down. The port is disabled if the configured threshold is exceeded. If no protocol option is specified, the threshold is applied to each of the tunneled Layer 2 protocol types. The range is 1 to 4096. The default is to have no threshold configured. Note If you also set a drop threshold on this interface, the shutdown-threshold value must be greater than or equal to the drop-threshold value.

	Command	Purpose
Step 6	<code>l2protocol-tunnel drop-threshold [cdp stp vtp] value</code>	(Optional) Configure the threshold in packets per second to be received for encapsulation before the interface drops packets. The port drops packets if the configured threshold is exceeded. If no protocol option is specified, the threshold is applied to each of the tunneled Layer 2 protocol types. The range is 1 to 4096. The default is to have no threshold configured. Note If you also set a shutdown threshold on this interface, the drop-threshold value must be less than or equal to the shutdown-threshold value.
Step 7	<code>exit</code>	Return to global configuration mode.
Step 8	<code>errdisable recovery cause l2ptguard</code>	(Optional) Configure the recovery mechanism from a Layer 2 maximum rate error so that the interface can be brought out of the disabled state and allowed to try again. You can also set the time interval. Errdisable recovery is disabled by default; when enabled, the default time interval is 300 seconds.
Step 9	<code>l2protocol-tunnel cos value</code>	(Optional) Configure the CoS value for all tunneled Layer 2 PDUs. The range is 0 to 7; the default is the default CoS value for the interface. If none is configured, the default is 5.
Step 10	<code>end</code>	Return to privileged EXEC mode.
Step 11	<code>show l2protocol</code>	Display the Layer 2 tunnel ports on the switch, including the protocols configured, the threshold, and the counters.
Step 12	<code>copy running-config startup-config</code>	(Optional) Save your entries in the configuration file.

Use the **`no l2protocol-tunnel [cdp | stp | vtp]`** interface configuration command to disable protocol tunneling for one of the Layer 2 protocols or for all three of them. Use the **`no l2protocol-tunnel shutdown-threshold [cdp | stp | vtp]`** and the **`no l2protocol-tunnel drop-threshold [cdp | stp | vtp]`** commands to return the shutdown and drop thresholds to the default settings.

This example shows how to configure Layer 2 protocol tunneling for CDP, STP, and VTP and to verify the configuration.

```
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# l2protocol-tunnel cdp
Switch(config-if)# l2protocol-tunnel stp
Switch(config-if)# l2protocol-tunnel vtp
Switch(config-if)# l2protocol-tunnel shutdown-threshold 1500
Switch(config-if)# l2protocol-tunnel drop-threshold 1000
Switch(config-if)# exit
Switch(config)# l2protocol-tunnel cos 7
Switch(config)# end
Switch# show l2protocol
COS for Encapsulated Packets: 7
```

Port	Protocol	Shutdown Threshold	Drop Threshold	Encapsulation Counter	Decapsulation Counter	Drop Counter
Gi1/0/2	cdp	1500	1000	0	0	0
	stp	1500	1000	0	0	0
	vtp	1500	1000	0	0	0

Monitoring and Maintaining Tunneling and Mapping Status

Table 13-2 shows the privileged EXEC commands for monitoring and maintaining 802.1Q and Layer 2 protocol tunneling and VLAN mapping.

Table 13-2 Commands for Monitoring and Maintaining Tunneling and VLAN Mapping

Command	Purpose
clear l2protocol-tunnel counters	Clear the protocol counters on Layer 2 protocol tunneling ports.
show dot1q-tunnel	Display 802.1Q tunnel ports on the switch.
show dot1q-tunnel interface <i>interface-id</i>	Verify if a specific interface is a tunnel port.
show l2protocol-tunnel	Display information about Layer 2 protocol tunneling ports.
show errdisable recovery	Verify if the recovery timer from a Layer 2 protocol-tunnel error disable state is enabled.
show interface <i>interface-id</i> vlan mapping	Display VLAN mapping information for the specified interface. The interface can be an ES port or a VLAN.
show l2protocol-tunnel interface <i>interface-id</i>	Display information about a specific Layer 2 protocol tunneling port.
show l2protocol-tunnel summary	Display only Layer 2 protocol summary information.
show vlan dot1q native	Display the status of native VLAN tagging on the switch.
show vlan mapping	Display VLAN mapping information (contents of the VLAN mapping table) for the ES ports.

For detailed information about these displays, refer to the command reference for this release.



Configuring STP

This chapter describes how to configure the Spanning Tree Protocol (STP) on port-based VLANs on the Catalyst 3750 Metro 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 15, “Configuring MSTP.”](#) For information about other spanning-tree features such as Port Fast, UplinkFast, root guard, and so forth, see [Chapter 16, “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, page 14-1](#)
- [Configuring Spanning-Tree Features, page 14-11](#)
- [Displaying the Spanning-Tree Status, page 14-22](#)

Understanding Spanning-Tree Features

These sections describe how basic spanning-tree features work:

- [STP Overview, page 14-2](#)
- [Spanning-Tree Topology and BPDUs, page 14-3](#)
- [Bridge ID, Switch Priority, and Extended System ID, page 14-4](#)
- [Spanning-Tree Interface States, page 14-4](#)
- [How a Switch or Port Becomes the Root Switch or Root Port, page 14-7](#)
- [Spanning Tree and Redundant Connectivity, page 14-8](#)
- [Spanning-Tree Address Management, page 14-8](#)
- [Accelerated Aging to Retain Connectivity, page 14-8](#)
- [Spanning-Tree Modes and Protocols, page 14-9](#)
- [Supported Spanning-Tree Instances, page 14-10](#)

- [Spanning-Tree Interoperability and Backward Compatibility, page 14-10](#)
- [STP and IEEE 802.1Q Trunks, page 14-10](#)
- [VLAN-Bridge Spanning Tree, page 14-11](#)

For configuration information, see the “Configuring Spanning-Tree Features” section on page 14-11.

For information about optional spanning-tree features, see [Chapter 16, “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 ports on a switch are part of a loop, the spanning-tree port priority and path cost settings control which port is put in the forwarding state and which is put in the blocking state. The spanning-tree port priority value represents the location of a port 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 controlled 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 14-1 on page 14-4](#).
- 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.

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 controls 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 switch supports the 802.1t spanning-tree extensions, and 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 14-1](#), 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.

Table 14-1 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.

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 14-14](#), the [“Configuring a Secondary Root Switch” section on page 14-16](#), and the [“Configuring the Switch Priority of a VLAN” section on page 14-19](#).

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 14-1 illustrates how an interface moves through the states.

Figure 14-1 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 switch interface. 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 interface moves to the listening state. An interface always enters the blocking state after switch initialization.

An interface in the blocking state performs these functions:

- Discards frames received on the interface
- 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 these functions:

- Discards frames received on the interface
- 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 these functions:

- Discards frames received on the interface
- 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 these functions:

- Receives and forwards frames received on the interface
- Forwards frames switched from another interface
- 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 these functions:

- Discards frames received on the interface
- 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 14-2](#), 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 14-2 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 10/100 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 port to a higher priority (lower numerical value) than the root port, the Gigabit Ethernet port 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, as shown in [Figure 14-3](#). Spanning tree automatically disables one interface but enables it if the other one fails. 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 14-3 Spanning Tree and Redundant Connectivity



You can also create redundant links between switches by using EtherChannel groups. For more information, see the [Chapter 27, “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, each switch receives but does not forward packets destined for addresses between 0x0180C2000000 and 0x0180C200000F.

If spanning tree is enabled, the CPU on the switch 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. It is the default spanning-tree mode used on all Ethernet, Fast Ethernet, and Gigabit Ethernet port-based VLANs. 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. 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 reprovision 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 15, “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 128 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 14-12](#).

Spanning-Tree Interoperability and Backward Compatibility

[Table 14-2](#) lists the interoperability and compatibility among the supported spanning-tree modes in a network.

Table 14-2 *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.

PVST+ is automatically enabled on 802.1Q trunks, and no user configuration is required. The external spanning-tree behavior on access ports and Inter-Switch Link (ISL) trunk ports is not affected by PVST+.

For more information on 802.1Q trunks, see [Chapter 10, “Configuring VLANs.”](#)

VLAN-Bridge Spanning Tree

Cisco VLAN-bridge spanning tree is used with the fallback bridging feature (bridge groups), which forwards non-IP protocols such as DECnet between two or more VLAN bridge domains or routed ports. The VLAN-bridge spanning tree allows the bridge groups to form a spanning tree on top of the individual VLAN spanning trees to prevent loops from forming if there are multiple connections among VLANs. It also prevents the individual spanning trees from the VLANs being bridged from collapsing into a single spanning tree.

To support VLAN-bridge spanning tree, some of the spanning-tree timers are increased. To use the fallback bridging feature, you must have the enhanced multilayer software image installed on your switch. For more information, see [Chapter 33, “Configuring Fallback Bridging.”](#)

Configuring Spanning-Tree Features

These sections describe how to configure spanning-tree features:

- [Default Spanning-Tree Configuration, page 14-11](#)
- [Spanning-Tree Configuration Guidelines, page 14-12](#)
- [Changing the Spanning-Tree Mode, page 14-13](#) (required)
- [Disabling Spanning Tree, page 14-14](#) (optional)
- [Configuring the Root Switch, page 14-14](#) (optional)
- [Configuring a Secondary Root Switch, page 14-16](#) (optional)
- [Configuring Port Priority, page 14-17](#) (optional)
- [Configuring Path Cost, page 14-18](#) (optional)
- [Configuring the Switch Priority of a VLAN, page 14-19](#) (optional)
- [Configuring Spanning-Tree Timers, page 14-20](#) (optional)

Default Spanning-Tree Configuration

[Table 14-3](#) shows the default spanning-tree configuration.

Table 14-3 Default Spanning-Tree Configuration

Feature	Default Setting
Enable state	Enabled on VLAN 1. For more information, see the “Supported Spanning-Tree Instances” section on page 14-10 .
Spanning-tree mode	PVST+. (Rapid PVST+ and MSTP are disabled.)
Switch priority	32768.
Spanning-tree port priority (configurable on a per-interface basis)	128.

Table 14-3 Default Spanning-Tree Configuration (continued)

Feature	Default Setting
Spanning-tree port cost (configurable on a per-interface basis)	1000 Mbps: 4. 100 Mbps: 19. 10 Mbps: 100.
Spanning-tree VLAN port priority (configurable on a per-VLAN basis)	128.
Spanning-tree VLAN port cost (configurable on a per-VLAN basis)	1000 Mbps: 4. 100 Mbps: 19. 10 Mbps: 100.
Spanning-tree timers	Hello time: 2 seconds. Forward-delay time: 15 seconds. Maximum-aging time: 20 seconds.

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 128 VLANs on the switch. The remaining VLANs operate with spanning tree disabled. However, you can map multiple VLANs to the same spanning-tree instances by using MSTP. For more information, see [Chapter 15, “Configuring MSTP.”](#)

If 128 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 control 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 14-10](#).

The UplinkFast and BackboneFast features are not supported with the rapid PVST+.

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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree mode {pvst mst rapid-pvst}	Configure a spanning-tree mode. • Select pvst to enable PVST+ (the default setting). • Select mst to enable MSTP (and RSTP). For more configuration steps, see Chapter 15, “Configuring MSTP.” • Select rapid-pvst to enable rapid PVST+.
Step 3	interface interface-id	(Recommended for rapid-PVST+ mode only) Specify an interface to configure, and enter interface configuration mode. Valid interfaces include physical ports, VLANs, and port channels. The VLAN ID range is 1 to 4094. The port-channel range is 1 to 12.
Step 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.
Step 5	end	Return to privileged EXEC mode.
Step 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 detects that this switch is running rapid PVST+.

	Command	Purpose
Step 7	show spanning-tree summary and show spanning-tree interface <i>interface-id</i>	Verify your entries.
Step 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 14-10. 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no spanning-tree vlan <i>vlan-id</i>	For <i>vlan-id</i> , the range is 1 to 4094.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
Step 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 software 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 14-1 on page 14-4](#).)

**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 through 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree vlan <i>vlan-id</i> root primary [diameter <i>net-diameter</i> [hello-time <i>seconds</i>]]	Configure a switch to become the root for the specified VLAN. - 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. The 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; the default is 2.
Step 3	end	Return to privileged EXEC mode.

	Command	Purpose
Step 4	show spanning-tree detail	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no spanning-tree vlan *vlan-id* root** global configuration command.

Configuring a Secondary Root Switch

When you configure the switch 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 that 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree vlan <i>vlan-id</i> root secondary [diameter <i>net-diameter</i> [hello-time <i>seconds</i>]]	Configure a switch to become the secondary root for the specified VLAN. - 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. The 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; the default is 2. 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 14-14.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree detail	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no spanning-tree vlan *vlan-id* root** global configuration command.

Configuring 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify an interface to configure, and enter interface configuration mode. Valid interfaces include physical ports and port-channel logical interfaces (port-channel <i>port-channel-number</i>).
Step 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. Valid values are 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224, and 240. All other values are rejected. The lower the number, the higher the priority.
Step 4	spanning-tree vlan <i>vlan-id</i> port-priority <i>priority</i>	Configure the port priority for a VLAN. - 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. The range is 1 to 4094. - For <i>priority</i>, the range is 0 to 240, in increments of 16; the default is 128. Valid values are 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224, and 240. All other values are rejected. The lower the number, the higher the priority.
Step 5	end	Return to privileged EXEC mode.
Step 6	show spanning-tree interface <i>interface-id</i> or show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
Step 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 to the 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 [“Configuring Trunk Ports for Load Sharing”](#) section on page 10-24.

Configuring 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 port 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify an interface to configure, and enter interface configuration mode. Valid interfaces include physical ports and port-channel logical interfaces (port-channel <i>port-channel-number</i>).
Step 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.
Step 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. - 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. The range is 1 to 4094. - For <i>cost</i>, the range is 1 to 200000000; the default value is derived from the media speed of the interface.
Step 5	end	Return to privileged EXEC mode.
Step 6	show spanning-tree interface <i>interface-id</i> or show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
Step 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 to the 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 [“Configuring Trunk Ports for Load Sharing”](#) section on page 10-24.

Configuring the Switch Priority of a VLAN

You can configure the switch priority and make it more likely that Catalyst 3750 Metro 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree vlan <i>vlan-id</i> priority <i>priority</i>	Configure the switch priority of a VLAN. - 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. The 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no spanning-tree vlan *vlan-id* priority** global configuration command.

Configuring Spanning-Tree Timers

Table 14-4 describes the timers that affect the entire spanning-tree performance.

Table 14-4 Spanning-Tree Timers

Variable	Description
Hello timer	Controls how often the switch broadcasts hello messages to other switches.
Forward-delay timer	Controls how long each of the listening and learning states last before the interface begins forwarding.
Maximum-age timer	Controls 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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. - 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. The range is 1 to 4094. - For <i>seconds</i>, the range is 1 to 10; the default is 2.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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 an interface waits before changing from its spanning-tree learning and listening states to the forwarding state. - 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. The range is 1 to 4094. - For <i>seconds</i>, the range is 4 to 30; the default is 15.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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. - 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. The range is 1 to 4094. - For <i>seconds</i>, the range is 6 to 40; the default is 20.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the 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 14-5](#):

Table 14-5 *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 interface 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.



Configuring MSTP

This chapter describes how to configure the Cisco implementation of the IEEE 802.1s Multiple STP (MSTP) on the Catalyst 3750 Metro 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 14, “Configuring STP.”](#) For information about other spanning-tree features such as Port Fast, UplinkFast, root guard, and so forth, see [Chapter 16, “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, page 15-2](#)
- [Understanding RSTP, page 15-6](#)
- [Configuring MSTP Features, page 15-11](#)
- [Displaying the MST Configuration and Status, page 15-23](#)

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, page 15-2](#)
- [IST, CIST, and CST, page 15-3](#)
- [Hop Count, page 15-5](#)
- [Boundary Ports, page 15-5](#)
- [“Interoperability with 802.1D STP” section on page 15-5](#)

For configuration information, see the [“Configuring MSTP Features” section on page 15-11](#).

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 15-1 on page 15-4](#).

The MST configuration controls 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 15-3](#) and the [“Operations Between MST Regions” section on page 15-4](#).

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 15-1 on page 15-4](#)), 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 15-1 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 15-1 MST Regions, IST Masters, and the CST Root




Figure 15-1 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 (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 also can 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. It cannot detect 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, page 15-6](#)
- [Rapid Convergence, page 15-7](#)
- [Synchronization of Port Roles, page 15-8](#)
- [Bridge Protocol Data Unit Format and Processing, page 15-9](#)

For configuration information, see the [“Configuring MSTP Features” section on page 15-11](#).

Port Roles and the Active Topology

The RSTP provides rapid convergence of the spanning tree by assigning port roles and by learning 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 14-3](#). 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 15-1 provides a comparison of 802.1D and RSTP port states.

Table 15-1 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 15-2, 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 learns 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 controlled by the duplex setting by using the **spanning-tree link-type** interface configuration command.

Figure 15-2 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 15-3](#).

Figure 15-3 Sequence of Events During Rapid Convergence

88761

Bridge Protocol Data Unit Format and Processing

The RSTP BPDU format is the same as the IEEE 802.1D BPDU format 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 15-2](#) shows the RSTP flag fields.

Table 15-2 RSTP BPDU 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 BPDU 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 BPDU 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) BPDU. 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, page 15-12](#)
- [MSTP Configuration Guidelines, page 15-12](#)
- [Specifying the MST Region Configuration and Enabling MSTP, page 15-13](#) (required)
- [Configuring the Root Switch, page 15-14](#) (optional)
- [Configuring a Secondary Root Switch, page 15-16](#) (optional)
- [Configuring Port Priority, page 15-17](#) (optional)
- [Configuring Path Cost, page 15-18](#) (optional)
- [Configuring the Switch Priority, page 15-19](#) (optional)
- [Configuring the Hello Time, page 15-19](#) (optional)
- [Configuring the Forwarding-Delay Time, page 15-20](#) (optional)
- [Configuring the Maximum-Aging Time, page 15-21](#) (optional)
- [Configuring the Maximum-Hop Count, page 15-21](#) (optional)
- [Specifying the Link Type to Ensure Rapid Transitions, page 15-22](#) (optional)
- [Restarting the Protocol Migration Process, page 15-22](#) (optional)

Default MSTP Configuration

Table 15-3 shows the default MSTP configuration.

Table 15-3 Default MSTP Configuration

Feature	Default Setting
Spanning-tree mode	PVST+ (Rapid PVST+ and MSTP are disabled).
Switch priority (configurable on a per-CIST port basis)	32768.
Spanning-tree port priority (configurable on a per-CIST port basis)	128.
Spanning-tree port cost (configurable on a per-CIST port basis)	1000 Mbps: 4. 100 Mbps: 19. 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 14-10.

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 14-10. For information on the recommended trunk port configuration, see the [“Interaction with Other Features”](#) section on page 10-20.
- 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree mst configuration	Enter MST configuration mode.
Step 3	instance <i>instance-id</i> vlan <i>vlan-range</i>	Map VLANs to an MST instance. • For <i>instance-id</i>, the range is 1 to 15. • For vlan <i>vlan-range</i>, the range is 1 to 4094. 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. To specify a VLAN range, use a hyphen; for example, instance 1 vlan 1-63 maps VLANs 1 through 63 to MST instance 1. 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.
Step 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.
Step 5	revision <i>version</i>	Specify the configuration revision number. The range is 0 to 65535.
Step 6	show pending	Verify your configuration by displaying the pending configuration.
Step 7	exit	Apply all changes, and return to global configuration mode.

	Command	Purpose
Step 8	spanning-tree mode mst	Enable MSTP. RSTP is also enabled.  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. You cannot run both MSTP and PVST+ or both MSTP and rapid PVST+ at the same time.
Step 9	end	Return to privileged EXEC mode.
Step 10	show running-config	Verify your entries.
Step 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. For a group of VLANs, the switch with the lowest bridge ID becomes the root switch.

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 14-1 on page 14-4](#).)

**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 through the **spanning-tree mst hello-time**, **spanning-tree mst forward-time**, and the **spanning-tree mst max-age** global configuration commands.

Beginning in privileged EXEC mode, follow these steps to configure a switch as the root switch. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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. - 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree mst <i>instance-id</i>	Verify your entries.
Step 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 the switch with the extended system ID support 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 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree mst instance-id root secondary [diameter <i>net-diameter</i> [hello-time <i>seconds</i>]]	Configure a switch as the secondary root switch. - 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. 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 15-14 .
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree mst instance-id	Verify your entries.
Step 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 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify an interface to configure, and enter interface configuration mode. Valid interfaces include physical ports and port-channel logical interfaces. The port-channel range is 1 to 12.
Step 3	spanning-tree mst <i>instance-id</i> port-priority <i>priority</i>	Configure the port priority. - 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. The 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show spanning-tree mst interface <i>interface-id</i> or show spanning-tree mst <i>instance-id</i>	Verify your entries.
Step 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 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify an interface to configure, and enter interface configuration mode. Valid interfaces include physical ports and port-channel logical interfaces. The port-channel range is 1 to 12.
Step 3	spanning-tree mst <i>instance-id</i> cost <i>cost</i>	Configure the cost. 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. - 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show spanning-tree mst interface <i>interface-id</i> or show spanning-tree mst <i>instance-id</i>	Verify your entries.
Step 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 interfaces 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree mst instance-id priority priority	Configure the switch priority. - 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. 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree mst instance-id	Verify your entries.
Step 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree mst	Verify your entries.
Step 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree mst	Verify your entries.
Step 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree mst	Verify your entries.
Step 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree mst	Verify your entries.
Step 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 15-7](#).

By default, the link type is controlled by the duplex mode of the port: 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 1	interface <i>interface-id</i>	Specify the interface to configure, and enter interface configuration mode. Valid interfaces include physical port, VLANs, and port-channel logical interfaces. The VLAN ID range is 1 to 4094. The port-channel range is 1 to 12.
Step 2	spanning-tree link-type point-to-point	Specify that the link type of a port is point-to-point.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree mst interface <i>interface-id</i>	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the port 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 also can 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 detect 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 port, 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 15-4](#):

Table 15-4 *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.

For information about other keywords for the **show spanning-tree** privileged EXEC command, refer to the command reference for this release.



Configuring Optional Spanning-Tree Features

This chapter describes how to configure optional spanning-tree features on the Catalyst 3750 Metro 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 14, “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 15, “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, page 16-1](#)
- [Configuring Optional Spanning-Tree Features, page 16-9](#)
- [Displaying the Spanning-Tree Status, page 16-16](#)

Understanding Optional Spanning-Tree Features

These sections describe how the optional spanning-tree features work:

- [Understanding Port Fast, page 16-2](#)
- [Understanding BPDU Guard, page 16-3](#)
- [Understanding BPDU Filtering, page 16-3](#)
- [Understanding UplinkFast, page 16-4](#)
- [Understanding BackboneFast, page 16-5](#)
- [Understanding Root Guard, page 16-7](#)
- [Understanding Loop Guard, page 16-8](#)

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 interfaces connected to a single workstation or server, as shown in Figure 16-1, to allow those devices to immediately connect to the network, rather than waiting for the spanning tree to converge.


Interfaces connected to a single workstation or server should not receive bridge protocol data units (BPDUs). An interface 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 interfaces must wait for spanning-tree to converge, it is effective only when used on interfaces connected to end stations. If you enable Port Fast on an interface 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 16-1 Port Fast-Enabled Interfaces



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 enable BPDU guard on Port Fast-enabled interfaces by using the **spanning-tree portfast bpduguard default** global configuration command. Spanning tree shuts down interfaces that are in a Port Fast-operational state. In a valid configuration, Port Fast-enabled interfaces do not receive BPDUs. Receiving a BPDU on a Port Fast-enabled interface signals an invalid configuration, such as the connection of an unauthorized device, and the BPDU guard feature puts the interface in the error-disabled state.

At the interface level, you enable BPDU guard on any interface by using the **spanning-tree bpduguard enable** interface configuration command without also enabling the Port Fast feature. When the interface 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 interface 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 interfaces by using the **spanning-tree portfast bpdupfilter default** global configuration command. This command prevents interfaces that are in a Port Fast-operational state from sending or receiving BPDUs. The interfaces 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 interfaces do not receive BPDUs. If a BPDU is received on a Port Fast-enabled interface, the interface loses its Port Fast-operational status, and BPDU filtering is disabled.

At the interface level, you can enable BPDU filtering on any interface by using the **spanning-tree bpdupfilter enable** interface configuration command without also enabling the Port Fast feature. This command prevents the interface 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 16-2 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 16-2 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 ports (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 16-3 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 16-3 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 interface on Switch C and transitions it to the forwarding state without going through the listening and learning states, as shown in Figure 16-4. This change takes approximately 1 to 5 seconds.

Figure 16-4 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 controls 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 interface 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 switch 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 interface, the root port and other blocked interfaces 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 interfaces become alternate paths to the root switch. If the inferior BPDU arrives on the root port and there are no blocked interfaces, 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 and waits for an RLQ reply from other switches in the network.


If the switch discovers that it still has an alternate path to the root, it expires the maximum aging time on the interface 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 interface that received the RLQ reply. If one or more alternate paths can still connect to the root switch, the switch makes all interfaces 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 16-5 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 16-5 BackboneFast Example Before Indirect Link Failure



If link L1 fails as shown in Figure 16-6, 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 interface on Switch C to move immediately to the listening state without waiting for the maximum aging time for the interface 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 16-6 shows how BackboneFast reconfigures the topology to account for the failure of link L1.

Figure 16-6 BackboneFast Example After Indirect Link Failure

If a new switch is introduced into a shared-medium topology as shown in [Figure 16-7](#), BackboneFast is not activated because the inferior BPDUs did not come from the recognized designated switch (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 switch to Switch A, the root switch.

Figure 16-7 Adding a Switch in a Shared-Medium Topology

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 16-8](#). 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 interface to be a designated port. If a boundary port is blocked in an internal spanning-tree (IST) instance because of root guard, the interface also is blocked in all MST instances. A boundary port is an interface 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 16-8 Root Guard in a Service-Provider Network



101232

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 enabled 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 interface is blocked by loop guard in all MST instances. On a boundary port, loop guard blocks the interface in all MST instances.

Configuring Optional Spanning-Tree Features

These sections describe how to configure optional spanning-tree features:

- [Default Optional Spanning-Tree Configuration, page 16-9](#)
- [Optional Spanning-Tree Configuration Guidelines, page 16-9](#)
- [Enabling Port Fast, page 16-10](#) (optional)
- [Enabling BPDU Guard, page 16-11](#) (optional)
- [Enabling BPDU Filtering, page 16-12](#) (optional)
- [Enabling UplinkFast for Use with Redundant Links, page 16-13](#) (optional)
- [Enabling BackboneFast, page 16-14](#) (optional)
- [Enabling Root Guard, page 16-14](#) (optional)
- [Enabling Loop Guard, page 16-15](#) (optional)

Default Optional Spanning-Tree Configuration

[Table 16-1](#) shows the default optional spanning-tree configuration.

Table 16-1 *Default Optional Spanning-Tree Configuration*

Feature	Default Setting
Port Fast, BPDU filtering, BPDU guard	Globally disabled (unless they are individually configured per interface).
UplinkFast	Globally disabled.
BackboneFast	Globally disabled.
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

An interface with the Port Fast feature enabled moves 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 an interface 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. For more information, see [Chapter 12, “Configuring Voice VLAN.”](#)

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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify an interface to configure, and enter interface configuration mode.
Step 3	spanning-tree portfast [trunk]	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.  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. By default, Port Fast is disabled on all interfaces.
Step 4	end	Return to privileged EXEC mode.
Step 5	show spanning-tree interface <i>interface-id</i> portfast	Verify your entries.
Step 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 interfaces that are Port Fast-enabled (the interfaces are in a Port Fast-operational state), spanning tree shuts down Port Fast-enabled interfaces that receive BPDUs.

In a valid configuration, Port Fast-enabled interfaces do not receive BPDUs. Receiving a BPDU on a Port Fast-enabled interface signals an invalid configuration, such as the connection of an unauthorized device, and the BPDU guard feature puts the interface in the error-disabled state. The BPDU guard feature provides a secure response to invalid configurations because you must manually put the interface 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 interfaces that connect to end stations; otherwise, an accidental topology loop could cause a data packet loop and disrupt switch and network operation.

You also can use the **spanning-tree bpduguard enable** interface configuration command to enable BPDU guard on any interface without also enabling the Port Fast feature. When the interface 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree portfast bpduguard default	Globally enable BPDU guard. By default, BPDU guard is disabled.
Step 3	interface <i>interface-id</i>	Specify the interface connected to an end station, and enter interface configuration mode.
Step 4	spanning-tree portfast	Enable the Port Fast feature.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 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 interfaces, it prevents interfaces that are in a Port Fast-operational state from sending or receiving BPDUs. The interfaces 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 interfaces do not receive BPDUs. If a BPDU is received on a Port Fast-enabled interface, the interface loses its Port Fast-operational status, and BPDU filtering is disabled.



Caution

Configure Port Fast only on interfaces 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 interface without also enabling the Port Fast feature. This command prevents the interface 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree portfast bpdudfilter default	Globally enable BPDU filtering. By default, BPDU filtering is disabled.
Step 3	interface <i>interface-id</i>	Specify the interface connected to an end station, and enter interface configuration mode.
Step 4	spanning-tree portfast	Enable the Port Fast feature.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable BPDU filtering, use the **no spanning-tree portfast bpdudfilter default** global configuration command.

You can override the setting of the **no spanning-tree portfast bpdudfilter default** global configuration command by using the **spanning-tree bpdudfilter enable** interface configuration command.

Enabling UplinkFast for Use with Redundant Links

UplinkFast cannot be enabled on VLANs that have been configured with a 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree summary	Verify your entries.
Step 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 reduce the chance that a 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree backbonefast	Enable BackboneFast.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree summary	Verify your entries.
Step 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 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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify an interface to configure, and enter interface configuration mode.
Step 3	spanning-tree guard root	Enable root guard on the interface. By default, root guard is disabled on all interfaces.

	Command	Purpose
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 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 interfaces 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.

	Command	Purpose
Step 1	show spanning-tree active	Verify which interfaces are alternate or root ports.
	or show spanning-tree mst	
Step 2	configure terminal	Enter global configuration mode.
Step 3	spanning-tree loopguard default	Enable loop guard. By default, loop guard is disabled.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 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 16-2](#):

Table 16-2 *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 interface 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.



Configuring IGMP Snooping and MVR

This chapter describes how to configure Internet Group Management Protocol (IGMP) snooping on the Catalyst 3750 Metro 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 IOS Release Network Protocols Command Reference, Part 1, for Release 12.1*.

This chapter consists of these sections:

- [Understanding IGMP Snooping, page 17-2](#)
- [Configuring IGMP Snooping, page 17-5](#)
- [Displaying IGMP Snooping Information, page 17-10](#)
- [Understanding Multicast VLAN Registration, page 17-12](#)
- [Configuring MVR, page 17-14](#)
- [Displaying MVR Information, page 17-18](#)
- [Configuring IGMP Filtering, page 17-19](#)
- [Displaying IGMP Filtering Configuration, page 17-22](#)



Note

You can either manage IP multicast group addresses through features such as IGMP snooping and MVR, or you can use static IP addresses.

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 (which could be a Catalyst 3750 Metro switch) sends out periodic general queries to all VLANs. All hosts interested in this multicast traffic send join requests and are added to the forwarding table entry. The switch creates one entry per VLAN in the IGMP snooping IP multicast forwarding table for each group from which it receives an IGMP join request.

The switch supports IP multicast group-based bridging, rather than MAC-addressed based groups. With multicast MAC address-based groups, if an IP address being configured translates (aliases) to a previously configured MAC address or to any reserved multicast MAC addresses (in the range 224.0.0.xxx), the command fails. Because the switch uses IP multicast groups, there are no address aliasing issues.

The IP multicast groups learned through IGMP snooping are dynamic. However, you can statically configure multicast groups by using the **ip igmp snooping vlan *vlan-id* static *ip_address* interface *interface-id*** 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.

These sections describe characteristics of IGMP snooping on the switch:

- [Joining a Multicast Group, page 17-2](#)
- [Leaving a Multicast Group, page 17-4](#)
- [Immediate-Leave Processing, page 17-5](#)

Joining a Multicast Group

When a host connected to the switch wants to join an IP multicast group, if it is an IGMP version 2 client, 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. IGMP version 1 or version 2 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 17-1](#).

Figure 17-1 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. 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 17-1](#), that includes the port numbers connected to Host 1 and the router.

Table 17-1 IGMP Snooping Forwarding Table

Destination Address	Type of Packet	Ports
224.1.2.3	IGMP	1, 2

The switch hardware can distinguish IGMP information packets from other packets for the multicast group. The information in the table tells the switching engine to send frames addressed to the 224.1.2.3 multicast IP address that are not IGMP packets 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 17-2](#)), the CPU receives that message and adds the port number of Host 4 to the forwarding table as shown in [Table 17-2](#). 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.

Figure 17-2 Second Host Joining a Multicast Group**Table 17-2 Updated IGMP Snooping Forwarding Table**

Destination Address	Type of Packet	Ports
224.1.2.3	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 only to those hosts listed in the forwarding table for that IP multicast group maintained by IGMP snooping.

When hosts want to leave a multicast group, they can 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 only supported with IGMP version 2 hosts.

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, page 17-5](#)
- [Enabling or Disabling IGMP Snooping, page 17-6](#)
- [Setting the Snooping Method, page 17-6](#)
- [Configuring a Multicast Router Port, page 17-8](#)
- [Configuring a Host Statically to Join a Group, page 17-9](#)
- [Enabling IGMP Immediate-Leave Processing, page 17-10](#)

Default IGMP Snooping Configuration

[Table 17-3](#) shows the default IGMP snooping configuration.

Table 17-3 *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

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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip igmp snooping	Globally enable IGMP snooping in all existing VLAN interfaces.
Step 3	end	Return to privileged EXEC mode.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip igmp snooping vlan <i>vlan-id</i>	Enable IGMP snooping on the VLAN interface. The VLAN ID range is 1 to 4094. Note IGMP snooping must be globally enabled before you can enable VLAN snooping.
Step 3	end	Return to privileged EXEC mode.
Step 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.

**Note**

If you want to use CGMP as the learning method and no multicast routers in the VLAN are CGMP proxy-enabled, you must enter the **ip cgmp router-only** command to dynamically access the router. For more information, see [Chapter 31, “Configuring IP Multicast Routing.”](#)

Beginning in privileged EXEC mode, follow these steps to alter the method in which a VLAN interface dynamically accesses a multicast router:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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: • 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip igmp snooping	Verify the configuration.
Step 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 and verify the configuration:

```
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.



Note

Static connections to multicast routers are supported only on switch ports.

Beginning in privileged EXEC mode, follow these steps to enable a static connection to a multicast router:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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. - The VLAN ID range is 1 to 4094. - The interface can be a physical interface or a port channel. The port channel range is 1 to 12.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip igmp snooping mrouter [vlan <i>vlan-id</i>]	Verify that IGMP snooping is enabled on the VLAN interface.
Step 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 gigabitethernet1/0/2
Switch(config)# end
Switch# show ip igmp snooping mrouter vlan 200
Vlan          ports
-----+-----
200           Gi1/0/2(static)nnnn
```

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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode
Step 2	ip igmp snooping vlan <i>vlan-id</i> static <i>ip_address</i> interface <i>interface-id</i>	Statically configure a Layer 2 port as a member of a multicast group: • <i>vlan-id</i> is the multicast group VLAN ID. • <i>ip_address</i> is the group IP address. • <i>interface-id</i> is the member port. It can be a physical interface or port channel (1 to 12).
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip igmp snooping multicast	Verify the member port and the IP address.
Step 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 224.1.2.3 interface gigabitethernet1/0/1
Switch(config)# end
```

```
Switch# show ip igmp snooping multicast
```

Vlan	Group Address	Type	Ports
----	-----	----	----
1	224.1.2.3	USER	Gi1/0/1

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.



Note

Immediate Leave is supported with only IGMP version 2 hosts.

Beginning in privileged EXEC mode, follow these steps to enable IGMP Immediate-Leave processing:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode
Step 2	ip igmp snooping vlan <i>vlan-id</i> immediate-leave	Enable IGMP Immediate-Leave processing on the VLAN interface.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip igmp snooping vlan <i>vlan-id</i>	Verify that Immediate Leave is enabled on the VLAN.
Step 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:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 130 immediate-leave
Switch(config)# end
```

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 17-4](#).

Table 17-4 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 multicast [count dynamic [count group <i>ip_address</i>] group <i>ip_address</i> user [count group <i>ip_address</i>]]	Display multicast table information for the switch or about a specific parameter: • count—Display the total number of entries for the specified command options instead of the actual entries. • dynamic—Display entries learned through IGMP snooping. • group <i>ip_address</i>—Display characteristics of the multicast group with the specified group IP address. • user—Display only the user-configured multicast entries.
show ip igmp snooping multicast vlan <i>vlan-id</i> [count dynamic [count group <i>ip_address</i>] group <i>ip_address</i> user [count group <i>ip_address</i>]]	Display multicast table information for a multicast VLAN or about a specific parameter for the VLAN: • count—Display the total number of entries for the specified command options instead of the actual entries. • dynamic—Display entries learned through IGMP snooping. • group <i>ip_address</i>—Display characteristics of the multicast group with the specified group IP address. • user—Display only the user-configured multicast entries.
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.

For more information about the keywords and options in these commands, refer to the command reference for this release.

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 IP multicast group 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.

You can set the switch for compatible or dynamic mode of MVR operation.


- In compatible mode, multicast data received by MVR hosts is forwarded to all MVR data ports, regardless of MVR host membership on those ports. The multicast data is forwarded only to those receiver ports which MVR hosts have explicitly joined, either by IGMP reports or by MVR static configuration. Also, IGMP reports received from MVR hosts are never forwarded out of MVR data ports that were configured in the switch.
- In dynamic mode, multicast data received by MVR hosts on the switch is forwarded from only those MVR data and client ports that the MVR hosts have explicitly joined, either by IGMP reports or by MVR static configuration. Any IGMP reports received from MVR hosts are also forwarded from all the MVR data ports in the switch. This eliminates using unnecessary bandwidth on MVR data port links, which occurs when the switch runs in compatible mode.

Only Layer 2 ports take part in MVR. You must configure ports as MVR receiver ports. Only one MVR multicast VLAN per switch is supported.

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. [Figure 17-3](#) is an example configuration. DHCP assigns an IP address to the set-top box or the PC. When a subscriber selects a channel, the set-top box or PC sends an IGMP report to Switch A to join the appropriate multicast. If the IGMP report matches one of the configured IP multicast group 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.

Figure 17-3 Multicast VLAN Registration Example



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 a MAC-based general 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 specified in the query. 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.

MVR eliminates the need to duplicate television-channel multicast traffic for subscribers in each VLAN. Multicast traffic for all channels is only sent around the VLAN trunk once—only on the multicast VLAN. The IGMP leave and join messages are in the VLAN to which the subscriber 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 (Switch A) modifies the forwarding behavior to allow the traffic to be forwarded from the multicast VLAN to the subscriber port in a different VLAN, selectively allowing traffic to cross between two VLANs.

IGMP reports are sent to the same IP multicast group address as the multicast data. The Switch A CPU must capture all IGMP join and leave messages from receiver ports and forward them to the multicast VLAN of the source (uplink) port, based on the MVR mode.

Configuring MVR

These sections include basic MVR configuration information:

- [Default MVR Configuration, page 17-14](#)
- [MVR Configuration Guidelines and Limitations, page 17-15](#)
- [Configuring MVR Global Parameters, page 17-15](#)
- [Configuring MVR Interfaces, page 17-17](#)

Default MVR Configuration

[Table 17-5](#) shows the default MVR configuration.

Table 17-5 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

Table 17-5 Default MVR Configuration (continued)

Feature	Default Setting
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 can only be access ports; they 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 (MVR group addresses) that can be configured on a switch (that is, the maximum number of television channels that can be received) is 256.
- MVR multicast data received in the source VLAN and leaving from receiver ports has its time-to-live (TTL) decremented by 1 in the switch.
- Because MVR on the switch uses IP multicast addresses instead of MAC multicast addresses, aliased IP multicast addresses are allowed on the switch. However, if the switch is interoperating with Catalyst 3550 or Catalyst 3500 XL switches, you should not configure IP addresses that alias between themselves or with the reserved IP multicast addresses (in the range 224.0.0.xxx).
- MVR is not supported when multicast routing is enabled on a switch. If you enable multicast routing and a multicast routing protocol while MVR is enabled, MVR is disabled, and you receive a warning message. If you try to enable MVR while multicast routing and a multicast routing protocol are enabled, the operation to enable MVR is cancelled, and you receive an error message.
- MVR can coexist with IGMP snooping on a switch.
- MVR data received on an MVR receiver port is not forwarded to MVR source ports.

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.



Note

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

Beginning in privileged EXEC mode, follow these steps to configure MVR parameters:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mvr	Enable MVR on the switch.

	Command	Purpose
Step 3	mvr group <i>ip-address</i> [<i>count</i>]	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.
Step 4	mvr querytime <i>value</i>	(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 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.
Step 6	mvr mode { dynamic compatible }	(Optional) Specify the MVR mode of operation: • 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.
Step 7	end	Return to privileged EXEC mode.
Step 8	show mvr or show mvr members	Verify the configuration.
Step 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 group address, set the query time to 1 second (10 tenths), specify the MVR multicast VLAN as VLAN 22, and set the MVR mode as dynamic:

```
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
```

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 Layer 2 MVR interfaces:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mvr	Enable MVR on the switch.
Step 3	interface <i>interface-id</i>	Enter interface configuration mode, and enter the type and number of the Layer 2 port to configure.
Step 4	mvr type {source receiver}	Configure an MVR port as one of these: • 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. The default configuration is as a non-MVR port. If you attempt to configure a non-MVR port with MVR characteristics, the operation fails.
Step 5	mvr vlan <i>vlan-id</i> group [<i>ip-address</i>]	(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. Note In compatible mode, this command applies to only receiver ports. In dynamic mode, it applies to receiver ports and source ports. Receiver ports can also dynamically join multicast groups by using IGMP join and leave messages.
Step 6	mvr immediate	(Optional) Enable the Immediate Leave feature of MVR on the port. Note This command applies to only receiver ports and should only be enabled on receiver ports to which a single receiver device is connected.
Step 7	end	Return to privileged EXEC mode.
Step 8	show mvr [interface members]	Verify the configuration.
Step 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 an interface 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 gigabitethernet1/0/2
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
Port      Type      Status      Immediate Leave
----      -
Gi1/0/2  RECEIVER  ACTIVE/DOWN  ENABLED
```

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 17-6](#) to display MVR configuration:

Table 17-6 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 vlan-id]</i>]	Displays all MVR interfaces and their MVR configurations. When a specific interface is entered, displays this information: - Type—Receiver or Source - Status—One of these: - 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; do not enter leading zeros.
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.

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.

IGMP filtering is only applicable to dynamic learning of IP multicast group addresses; not static configuration.

You can also set the maximum number of IGMP groups that a Layer 2 interface can join.

Default IGMP Filtering Configuration

Table 17-7 shows the default IGMP filtering configuration.

Table 17-7 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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 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.
Step 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.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip igmp profile <i>profile number</i>	Verify the profile configuration.
Step 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 access ports only; you cannot apply IGMP profiles to routed ports or SVIs. 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the physical interface to configure. The interface must be a Layer 2 port that does not belong to an EtherChannel port group.
Step 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 4	end	Return to privileged EXEC mode.
Step 5	show running-config interface <i>interface-id</i>	Verify the configuration.
Step 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 a port:

```
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# ip igmp filter 4
Switch(config-if)# 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.

This restriction can be applied to Layer 2 ports only; you cannot set a maximum number of IGMP groups on routed ports or SVIs. You also 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the physical interface to configure. The interface must be a Layer 2 port that does not belong to an EtherChannel group.
Step 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config interface <i>interface-id</i>	Verify the configuration.
Step 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 a port can join to 25.

```
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# ip igmp max-groups 25
Switch(config-if)# 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 17-8](#) to display IGMP filtering configuration:

Table 17-8 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-config [interface interface-id]	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.



Configuring Port-Based Traffic Control

This chapter describes how to configure port-based traffic control features on the Catalyst 3750 Metro 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, page 18-1](#)
- [Configuring Protected Ports, page 18-4](#)
- [Configuring Port Blocking, page 18-5](#)
- [Configuring Port Security, page 18-6](#)
- [Displaying Port-Based Traffic Control Settings, page 18-14](#)

Configuring Storm Control

These sections include storm control configuration information and procedures:

- [Understanding Storm Control, page 18-1](#)
- [Default Storm Control Configuration, page 18-3](#)
- [Enabling Storm Control, page 18-3](#)

Understanding Storm Control

Storm control prevents switchports on a LAN from being disrupted by a broadcast, multicast, or unicast storm on one of the physical interfaces. A LAN storm occurs when packets flood the LAN, creating excessive traffic and degrading network performance. Errors in the protocol-stack implementation or in the network configuration can cause a storm.

Storm control (or traffic suppression) monitors incoming traffic statistics over a time period and compares the measurement with a predefined suppression level threshold. The threshold represents the percentage of the total available bandwidth of the port. The switch supports separate storm control thresholds for broadcast, multicast, and unicast traffic. If the threshold of a traffic type is reached, further traffic of that type is suppressed until the incoming traffic falls below the threshold level.


**Note**

When the storm control threshold for multicast traffic is reached, all multicast traffic except control traffic, such as bridge protocol data unit (BDPU) and Cisco Discovery Protocol (CDP) frames, are blocked. However, the switch does not differentiate between routing updates, such as OSPF, and regular multicast data traffic, so both types of traffic are blocked.

When storm control is enabled, the switch monitors packets passing from an interface to the switching bus and determines if the packet is unicast, multicast, or broadcast. The switch monitors the number of broadcast, multicast, or unicast packets received within a 200-millisecond time interval, and when a threshold for one type of traffic is reached, that type of traffic is dropped. This threshold is specified as a percentage of total available bandwidth that can be used by broadcast (multicast or unicast) traffic.

The graph in [Figure 18-1](#) shows broadcast traffic patterns on an interface over a given period of time. The example can also be applied to multicast and unicast traffic. In this example, the broadcast traffic being forwarded exceeded the configured threshold between time intervals T1 and T2 and between T4 and T5. When the amount of specified traffic exceeds the threshold, all traffic of that kind is dropped for the next time period. Therefore, broadcast traffic is blocked during the intervals following T2 and T5. At the next time interval (for example, T3), if broadcast traffic does not exceed the threshold, it is again forwarded.

Figure 18-1 Broadcast Storm Control Example



The combination of the storm-control suppression level and the 200-millisecond time interval control the way the storm control algorithm works. A higher threshold allows more packets to pass through. A threshold value of 100 percent means that no limit is placed on the traffic. A value of 0.0 means that all broadcast, multicast, or unicast traffic on that port is blocked.

**Note**

Because packets do not arrive at uniform intervals, the 200-millisecond time interval during which traffic activity is measured can affect the behavior of storm control.

The switch continues to monitor traffic on the port, and when the utilization level is below the threshold level, the type of traffic that was dropped is forwarded again.

You use the **storm-control** interface configuration commands to set the threshold value for each traffic type.

**Note**

Although visible in the command-line interface (CLI) online help, the **switchport broadcast**, **switchport multicast**, and **switchport unicast** interface configuration commands for setting suppression levels are not available. These commands are obsolete, replaced by the **storm-control** interface configuration commands.

Default Storm Control Configuration

By default, unicast, broadcast, and multicast storm control is disabled on the switch interfaces; that is, the suppression level is 100 percent.

Enabling Storm Control

You enable storm control on an interface and enter the percentage of total available bandwidth that you want to be used by a particular type of traffic; entering 100 percent allows all traffic. However, because of hardware limitations and the way in which packets of different sizes are counted, threshold percentages are approximations. Depending on the sizes of the packets making up the incoming traffic, the actual enforced threshold might differ from the configured level by several percentage points.

**Note**

Storm control is supported only on physical interfaces; it is not supported on EtherChannel port channels even though the command is available in the CLI.

Beginning in privileged EXEC mode, follow these steps to enable a particular type of storm control:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the type and number of the physical interface to configure, for example gigabitethernet1/0/1 .
Step 3	storm-control broadcast level <i>level</i> [<i>.level</i>]	Specify the broadcast traffic suppression level for an interface as a percentage of total bandwidth. The level can be from 1 to 100; the optional fraction of a level can be from 0 to 99. A threshold value of 100 percent means that no limit is placed on broadcast traffic. A value of 0.0 means that all broadcast traffic on that port is blocked.
Step 4	storm-control multicast level <i>level</i> [<i>.level</i>]	Specify the multicast traffic suppression level for an interface as a percentage of total bandwidth. The level can be from 1 to 100; the optional fraction of a level can be from 0 to 99. A threshold value of 100 percent means that no limit is placed on broadcast traffic. A value of 0.0 means that all multicast traffic on that port is blocked.
Step 5	storm-control unicast level <i>level</i> [<i>.level</i>]	Specify the unicast traffic suppression level for an interface as a percentage of total bandwidth. The level can be from 1 to 100; the optional fraction of a level can be from 0 to 99. A threshold value of 100 percent means that no limit is placed on broadcast traffic. A value of 0.0 means that all unicast traffic on that port is blocked.
Step 6	end	Return to privileged EXEC mode.

	Command	Purpose
Step 7	show storm-control [<i>interface-id</i>] [broadcast multicast unicast]	Verify the storm control suppression levels set on the interface for the specified traffic type. If you do not enter a traffic type, broadcast storm control settings are displayed.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable storm control, use the **no storm-control broadcast level**, **no storm-control multicast level**, or **no storm-control unicast level** interface configuration commands.

This example shows how to set the multicast storm control level at 70.5 percent on an interface and verify the configuration:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# storm-control multicast level 70.5
Switch(config-if)# end
Switch# show storm-control gigabitethernet1/0/17 multicast
Interface  Filter State    Level    Current
-----  -
Gi1/0/1    Forwarding    70.50%   0.00%
```

This example shows how to disable the multicast storm control on an interface and verify the configuration:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# no storm-control multicast level
Switch(config-if)# end
Switch# show storm-control gigabitethernet1/0/1 multicast
Interface  Filter State    Level    Current
-----  -
Gi1/0/1    inactive      100.00%  N/A
```

Configuring Protected Ports

Some applications require that no traffic be forwarded at Layer 2 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.

Default Protected Port Configuration

The default is to have no protected ports defined.

Protected Port Configuration Guidelines

You can configure protected ports on a physical interface (for example, Gigabit Ethernet1/0/1) 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.

Configuring a Protected Port

Beginning in privileged EXEC mode, follow these steps to define a port as a protected port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the type and number of the interface to configure, for example gigabitethernet1/0/1 .
Step 3	switchport protected	Configure the interface to be a protected port.
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces <i>interface-id</i> switchport	Verify your entries.
Step 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 an interface as a protected port:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# switchport protected
Switch(config-if)# end
```

Configuring Port Blocking

By default, the switch floods packets with unknown destination MAC addresses out of all ports. If unknown unicast and multicast traffic is forwarded to a protected port, there could be security issues. To prevent unknown unicast or multicast traffic from being forwarded from one port to another, you can block a port (protected or nonprotected) from flooding unknown unicast or multicast packets to other ports.

Default Port Blocking Configuration

The default is to not block flooding of unknown multicast and unicast traffic out of a port, but to flood these packets to all ports.

Blocking Flooded Traffic on an Interface



Note

The interface can be a physical interface (for example, Gigabit Ethernet 1/0/1) or an EtherChannel group (for example, port-channel 5). When you block multicast or unicast traffic for a port channel, it is blocked on all ports in the port channel group.

Beginning in privileged EXEC mode, follow these steps to disable the flooding of multicast and unicast packets out of an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the type and number of the interface to configure, for example gigabitethernet1/0/1 .
Step 3	switchport block multicast	Block unknown multicast forwarding out of the port.
Step 4	switchport block unicast	Block unknown unicast forwarding out of the port.
Step 5	end	Return to privileged EXEC mode.
Step 6	show interfaces <i>interface-id</i> switchport	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the interface to the default condition where no traffic is blocked and normal forwarding occurs on the port, use the **no switchport block {multicast | unicast}** interface configuration commands.

This example shows how to block unicast and multicast flooding on an interface:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# switchport block multicast
Switch(config-if)# switchport block unicast
Switch(config-if)# end
```

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. If you limit the number of secure MAC addresses to one and assign a single secure MAC address, the workstation attached to that port is assured the full bandwidth of the port.

If a port is configured as a secure port and the maximum number of secure MAC addresses is reached, when the MAC address of a station attempting to access the port is different from any of the identified secure MAC addresses, a security violation occurs. Also, if a station with a secure MAC address configured or learned on one secure port attempts to access another secure port, a violation is flagged.

These sections include port security configuration information and procedures:

- [Understanding Port Security, page 18-7](#)
- [Default Port Security Configuration, page 18-9](#)
- [Configuration Guidelines, page 18-9](#)
- [Enabling and Configuring Port Security, page 18-10](#)
- [Enabling and Configuring Port Security Aging, page 18-13](#)

Understanding Port Security

This section contains information about these topics:

- [Secure MAC Addresses, page 18-7](#)
- [Security Violations, page 18-8](#)

Secure MAC Addresses

You configure the maximum number of secure addresses allowed on a port by using the **switchport port-security maximum** *value* interface configuration command.



Note

If you try to set the maximum value to a number less than the number of secure addresses already configured on an interface, the command is rejected.

The switch supports 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 configured, 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, when the switch restarts, the interface does not need to dynamically reconfigure them.

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. All sticky secure MAC addresses are added to the running configuration.

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 sticky secure addresses, 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.

The maximum number of secure MAC addresses that you can configure on a switch is determined by the maximum number of available MAC addresses allowed in the system. This number is determined by the active Switch Database Management (SDM) template. See [Chapter 6, “Configuring SDM Templates.”](#) This number represents the total of available MAC addresses, including those used for other Layer 2 functions and any other secure MAC addresses configured on interfaces.

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 maximum limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses to drop below the maximum value or increase the number of maximum allowable addresses. You are not notified that a security violation has occurred.



Note

We do not recommend configuring the protect violation mode on a trunk port. The protect mode disables learning when any VLAN reaches its maximum limit, even if the port has not reached its maximum limit.

- **restrict**—when the number of secure MAC addresses reaches the maximum limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses to drop below the maximum value or increase the number of maximum allowable addresses. In this mode, you are notified that a security violation has occurred. An SNMP trap is sent, a syslog message is logged, and the violation counter increments.
- **shutdown**—a port security violation causes the interface to become error-disabled and to shut down immediately, and the port LED turns off. An SNMP trap is sent, a syslog message is logged, and the violation counter increments. 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 shut down** interface configuration commands. This is the default mode.

[Table 18-1](#) shows the violation mode and the actions taken when you configure an interface for port security.

Table 18-1 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 returns an error message if you manually configure an address that would cause a security violation.

Default Port Security Configuration

Table 18-2 shows the default port security configuration for an interface.

Table 18-2 Default Port Security Configuration

Feature	Default Setting
Port security	Disabled on a port.
Sticky address learning	Disabled.
Maximum secure MAC addresses per port	1.
Violation mode	Shutdown. The port shuts down when the maximum number of secure MAC addresses is exceeded.
Port security aging	Disabled. Aging time is 0. Static aging is disabled. Type is absolute.

Configuration Guidelines

Follow these guidelines when configuring port security:

- Port security can only be configured on static access ports or trunk ports. A secure port cannot be a dynamic access port.
- A secure port cannot be a destination port for Switched Port Analyzer (SPAN).
- A secure port cannot belong to a Fast EtherChannel or a Gigabit EtherChannel port group.
- You cannot configure static secure or sticky secure MAC addresses in the voice VLAN.



Note Voice VLAN is only supported on access ports and not on trunk ports, even though the configuration is allowed.

- 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 two plus the maximum number of secure addresses allowed on the access VLAN. When the port is connected to a Cisco IP phone, the IP phone requires up to two MAC addresses. The IP phone address is learned on the voice VLAN and might also be learned on the access VLAN. Connecting a PC to the IP phone requires additional MAC addresses.
- If any type of port security is enabled on the access VLAN, dynamic port security is automatically enabled on the voice VLAN. You cannot configure port security on a per-VLAN basis.
- When a voice VLAN is configured on a secure port that is also configured as a sticky secure port, all addresses 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.
- When you enter a maximum secure address value for an interface, and the new value is greater than the previous value, the new value overwrites the previously configured value. If the new value is less than the previous value and the number of configured secure addresses on the interface exceeds the new value, the command is rejected.
- The switch does not support port security aging of sticky secure MAC addresses.

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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the physical interface to configure, for example gigabitethernet1/0/1 .
Step 3	switchport mode { access trunk }	Set the interface switchport mode as access or trunk; an interface in the default mode (dynamic auto) cannot be configured as a secure port.
Step 4	switchport port-security	Enable port security on the interface.
Step 5	switchport port-security maximum <i>value</i> [vlan [<i>vlan-list</i>]]	(Optional) Set the maximum number of secure MAC addresses for the interface. The maximum number of secure MAC addresses that you can configure on a switch is determined by the maximum number of available MAC addresses allowed in the system. This number is determined by the active Switch Database Management (SDM) template. See Chapter 6, “Configuring SDM Templates.” This number represents the total of available MAC addresses, including those used for other Layer 2 functions and any other secure MAC addresses configured on interfaces. (Optional) For trunk ports, you can set the maximum number of secure MAC addresses on a VLAN. If the vlan keyword is not entered, the default value is used. • vlan—set a per-VLAN maximum value. • vlan <i>vlan-list</i>—set a per-VLAN maximum value on a range of VLANs separated by a hyphen, or a series of VLANs separated by commas. For non-specified VLANs, the per-VLAN maximum value is used.

	Command	Purpose
Step 6	switchport port-security violation {protect restrict shutdown}	(Optional) Set the violation mode, the action to be taken when a security violation is detected, as one of these: • protect—When the number of port secure MAC addresses reaches the maximum limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses to drop below the maximum value or increase the number of maximum allowable addresses. You are not notified that a security violation has occurred. Note We do not recommend configuring the protect mode on a trunk port. The protect mode disables learning when any VLAN reaches its maximum limit, even if the port has not reached its maximum limit. • 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. An SNMP trap is sent, a syslog message is logged, and the violation counter increments. • shutdown—The interface is error-disabled when a violation occurs, and the port LED turns off. An SNMP trap is sent, a syslog message is logged, and the violation counter increments. 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.
Step 7	switchport port-security mac-address <i>mac-address</i> [vlan <i>vlan-id</i>]	(Optional) Enter a secure MAC address for the interface. 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. (Optional) On a trunk port, you can specify the VLAN ID and the MAC address. If no VLAN ID is specified, the native VLAN is used. 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.
Step 8	switchport port-security mac-address sticky	(Optional) Enable sticky learning on the interface.
Step 9	switchport port-security mac-address sticky <i>mac-address</i>	(Optional) Enter a sticky secure MAC address, repeating the command as many times as necessary. If you configure fewer secure MAC addresses than the maximum, the remaining MAC addresses are dynamically learned, are converted to sticky secure MAC addresses, and are added to the running configuration. Note If you do not enable sticky learning before this command is entered, an error message appears, and you cannot enter a sticky secure MAC address.
Step 10	end	Return to privileged EXEC mode.

	Command	Purpose
Step 11	show port-security	Verify your entries.
Step 12	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 {protocol | 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. However, if you have previously saved the configuration with the sticky MAC addresses, you should save the configuration again after entering the **no switchport port-security mac-address sticky** command, or the sticky addresses will be restored if the switch reboots.

To delete a specific secure MAC address from the address table, use the **no switchport port-security mac-address mac-address** interface configuration command.

To delete all dynamic secure addresses on an interface from the address table, enter the **no switchport port-security** interface configuration command followed by the **switchport port-security** command (to re-enable port security on the interface). If you use the **no switchport port-security mac-address sticky** interface configuration command to convert sticky secure MAC addresses to dynamic secure MAC addresses before entering the **no switchport port-security** command, all secure addresses on the interface except those that were manually configured are deleted.

You must specifically delete configured secure MAC addresses from the address table by using the **no switchport port-security mac-address mac-address** interface configuration command.

This example shows how to enable port security on an interface 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(config)# interface gigabitethernet1/0/1
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
```

This example shows how to configure a static secure MAC address on VLAN 3 on an interface:

```
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# switchport mode trunk
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security mac-address 0000.02000.0004 vlan 3
```

Enabling and Configuring Port Security Aging

You can use port security aging to set the aging time for all 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 devices 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 secure addresses on a per-port basis.

Beginning in privileged EXEC mode, follow these steps to configure port security aging:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode for the port on which you want to enable port security aging.
Step 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. Note The switch does not support port security aging of sticky secure addresses. 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: • absolute—Sets the aging type as absolute aging. All the secure addresses on this port age out exactly after the time (minutes) specified lapses and are removed from the secure address list. • 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show port-security [interface <i>interface-id</i>] [address]	Verify your entries.
Step 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 an interface:

```
Switch(config)# interface gigabitethernet1/0/1
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 18-3](#).

Table 18-3 Commands for Displaying Traffic Control Status and Configuration

Command	Purpose
show interfaces [<i>interface-id</i>] switchport	Displays the administrative and operational status of all switching (nonrouting) ports or the specified port, including port blocking and port protection settings.
show storm-control [<i>interface-id</i>] [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 [<i>interface-id</i>] 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 [<i>interface-id</i>] 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 [<i>interface-id</i>] 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 <i>interface-id</i>]	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 <i>interface-id</i>] address	Displays all secure MAC addresses configured on all switch interfaces or on a specified interface with aging information for each address.
show port-security interface <i>interface-id</i> vlan	Displays the number of secure MAC addresses configured per VLAN on the specified interface.



Configuring CDP

This chapter describes how to configure Cisco Discovery Protocol (CDP) on the Catalyst 3750 Metro switch.



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 Configuration Fundamentals Command Reference for Release 12.1*.

This chapter consists of these sections:

- [Understanding CDP, page 19-1](#)
- [Configuring CDP, page 19-2](#)
- [Monitoring and Maintaining CDP, page 19-5](#)

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.

The switch supports CDP version 2.

Configuring CDP

These sections include CDP configuration information and procedures:

- [Default CDP Configuration, page 19-2](#)
- [Configuring the CDP Characteristics, page 19-2](#)
- [Disabling and Enabling CDP, page 19-3](#)
- [Disabling and Enabling CDP on an Interface, page 19-4](#)

Default CDP Configuration

[Table 19-1](#) shows the default CDP configuration.

Table 19-1 *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.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	cdp timer <i>seconds</i>	(Optional) Set the transmission frequency of CDP updates in seconds. The range is 5 to 254; the default is 60 seconds.
Step 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 10 to 255 seconds; the default is 180 seconds.
Step 4	cdp advertise-v2	(Optional) Configure CDP to send version-2 advertisements. This is the default state.
Step 5	end	Return to privileged EXEC mode.

	Command	Purpose
Step 6	show cdp	Verify configuration by displaying global information about CDP on the device.
Step 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
  Sending CDPv2 advertisements is enabled
```

For additional CDP **show** commands, see the [“Monitoring and Maintaining CDP” section on page 19-5](#).

Disabling and Enabling CDP

CDP is enabled by default.

Beginning in privileged EXEC mode, follow these steps to disable the CDP device discovery capability:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no cdp run	Disable CDP.
Step 3	end	Return to privileged EXEC mode.

Beginning in privileged EXEC mode, follow these steps to enable CDP when it has been disabled:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	cdp run	Enable CDP after disabling it.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the interface on which you are disabling CDP.
Step 3	no cdp enable	Disable CDP on an interface.
Step 4	end	Return to privileged EXEC mode.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the interface on which you are enabling CDP.
Step 3	cdp enable	Enable CDP on an interface after disabling it.
Step 4	end	Return to privileged EXEC mode.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to enable CDP on a port when it has been disabled.

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/0/1
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>interface-id</i>]	Display information about interfaces where CDP is enabled. Enter an interface ID to display CDP information for that interface only.
show cdp neighbors [<i>interface-id</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 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
```




Configuring UDLD

This chapter describes how to configure the UniDirectional Link Detection (UDLD) protocol on the Catalyst 3750 Metro 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 UDLD, page 20-1](#)
- [Configuring UDLD, page 20-4](#)
- [Displaying UDLD Status, page 20-6](#)

Understanding UDLD

UDLD is a Layer 2 protocol that enables devices connected through fiber-optic or twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. 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 interface and alerts you. Unidirectional links can cause a variety of problems, including spanning-tree topology loops.

Modes of Operation

UDLD supports two modes of operation: normal (the default) and aggressive. In normal mode, UDLD can detect unidirectional links due to misconnected ports on fiber-optic connections. In aggressive mode, UDLD also can detect unidirectional links due to one-way traffic on fiber-optic and on twisted-pair links and to misconnected ports on fiber-optic links.

In normal and in aggressive mode, UDLD works with the Layer 1 mechanisms to learn 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, the 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 a local device is received by its neighbor but traffic from the neighbor is not received by the local device.

In normal mode, UDLD can detect a unidirectional link when fiber strands in a fiber-optic port are misconnected and the Layer 1 mechanisms do not detect this misconnection. If the ports are connected correctly but the traffic is one way, UDLD does not detect the unidirectional link because the Layer 1 mechanism, which is supposed to detect this condition, does not do so. In this case, the logical link is considered undetermined, and UDLD does not disable the interface.

When UDLD is in normal mode, if one of the fiber strands in a pair is disconnected, as long as autonegotiation is active, the link does not stay up because the Layer 1 mechanisms did not detect a physical problem with the link. In this case, UDLD does not take any action and the logical link is considered undetermined.

In aggressive mode, UDLD can detect a unidirectional link by using the previous detection methods. UDLD in aggressive mode also can detect a unidirectional link on a point-to-point link on which no failure between the two devices is allowed. It also can detect a unidirectional link when one of these problems exists:

- On fiber-optic or twisted-pair links, one of the interfaces cannot send or receive traffic.
- On fiber-optic or twisted-pair links, one of the interfaces is down while the other is up.
- One of the fiber strands in a pair is disconnected.

In these cases, UDLD shuts down the affected interface.

In a point-to-point link, UDLD hello packets can be considered as a heart beat whose presence guarantees the health of the link. Conversely, the loss of the heart beat means that the link must be shut down if it is not possible to re-establish a bidirectional link.

If both fibers are working normally from a Layer 1 perspective, UDLD in aggressive mode detects whether those fibers 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.

Methods to Detect Unidirectional Links

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 might shut down, depending on the UDLD mode. When UDLD is in normal mode, the link might be considered undetermined and might not shut down. When UDLD is in aggressive mode, the link is considered unidirectional, and the interface is shut down.

If UDLD in normal mode is in the advertisement or in the detection phase and all the neighbor cache entries are aged out, UDLD restarts the link-up sequence to resynchronize with any potentially out-of-sync neighbors.

If you enable aggressive mode when all the neighbors of an interface have aged out either in the advertisement or in the detection phase, UDLD restarts the link-up sequence to resynchronize with any potentially out-of-sync neighbor. UDLD shuts down the interface if, after the fast train of messages, the link state is still undetermined.

Figure 20-1 shows an example of a unidirectional link condition.

Figure 20-1 UDLD Detection of a Unidirectional Link



98648

Configuring UDLD

This section describes how to configure UDLD on your switch. It contains this configuration information:

- [Default UDLD Configuration, page 20-4](#)
- [UDLD Configuration Guidelines, page 20-4](#)
- [Enabling UDLD Globally, page 20-5](#)
- [Enabling UDLD on an Interface, page 20-5](#)
- [Resetting an Interface Disabled by UDLD, page 20-6](#)

Default UDLD Configuration

[Table 20-1](#) shows the default UDLD configuration.

Table 20-1 Default UDLD Configuration

Feature	Default Setting
UDLD global enable state	Globally disabled
UDLD per-port enable state for fiber-optic media	Disabled on all Ethernet fiber-optic ports
UDLD per-port enable state for twisted-pair (copper) media	Disabled on all Ethernet 10/100 and 1000BASE-TX ports
UDLD aggressive mode	Disabled

UDLD Configuration Guidelines

These are the UDLD configuration guidelines:

- UDLD is not supported on ATM ports.
- A UDLD-capable interface also cannot detect a unidirectional link if it is connected to a UDLD-incapable interface of another switch.
- When configuring the mode (normal or aggressive), make sure that the same mode is configured on both sides of the link.

Enabling UDLD Globally

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 ports on the switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	udld {aggressive enable message time message-timer-interval}	Specify the UDLD mode of operation: • aggressive—Enables UDLD in aggressive mode on all fiber-optic ports. • enable—Enables UDLD in normal mode on all fiber-optic ports on the switch. UDLD is disabled by default. An individual interface configuration overrides the setting of the udld enable global configuration command. For more information about aggressive and normal modes, see the “Modes of Operation” section on page 20-1. • message time message-timer-interval—Configures the period of time between UDLD probe messages on interfaces that are in the advertisement phase and are detected to be bidirectional. The range is 7 to 90 seconds. Note This command affects fiber-optic ports 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 20-5.
Step 3	end	Return to privileged EXEC mode.
Step 4	show udld	Verify your entries.
Step 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 either to enable UDLD in the aggressive or normal mode or to disable UDLD on an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface interface-id	Specify the interface to be enabled for UDLD, and enter interface configuration mode.

	Command	Purpose
Step 3	udld port { aggressive disable }	UDLD is disabled by default. • udld port—Enables UDLD in normal mode on the specified interface. • udld port aggressive—Enables UDLD in aggressive mode on the specified interface. • udld port disable—Disables UDLD on the specified fiber-optic port. This command overrides the UDLD global setting and is available only on fiber-optic ports. For more information about aggressive and normal modes, see the “Modes of Operation” section on page 20-1 .
Step 4	end	Return to privileged EXEC mode.
Step 5	show udld <i>interface-id</i>	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Resetting an Interface Disabled by UDLD

Beginning in privileged EXEC mode, follow these steps to reset all interfaces disabled by UDLD:

	Command	Purpose
Step 1	udld reset	Reset all interfaces disabled by UDLD.
Step 2	show udld	Verify your entries.

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 **no udld {aggressive | enable}** global configuration command followed by the **udld {aggressive | enable}** global configuration command re-enables the disabled interfaces.
- The **udld port disable** interface configuration command followed by the **udld port [aggressive]** interface configuration command re-enables the disabled fiber-optic port.
- 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** 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.



Configuring SPAN and RSPAN

This chapter describes how to configure Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) on the Catalyst 3750 Metro 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, page 21-1](#)
- [Configuring SPAN and RSPAN, page 21-9](#)
- [Displaying SPAN and RSPAN Status, page 21-23](#)

Understanding SPAN and RSPAN

You can analyze network traffic passing through ports or VLANs by using SPAN or RSPAN to send a copy of the traffic to another port on the switch or on another switch that has been connected to a network analyzer or other monitoring or security device. SPAN copies (or mirrors) traffic received or sent (or both) on source ports or source VLANs to a destination port for analysis. SPAN does not affect the switching of network traffic on the source ports or VLANs. You must dedicate the destination port for SPAN use. Except for traffic that is required for the SPAN or RSPAN session, destination ports do not receive or forward traffic.

Only traffic that enters or leaves source ports or traffic that enters or leaves source VLANs can be monitored by using SPAN; traffic routed to a source VLAN cannot be monitored. For example, if incoming traffic is being monitored, traffic that gets routed from another VLAN to the source VLAN cannot be monitored; however, traffic that is received on the source VLAN and routed to another VLAN can be monitored.

Enhanced-services (ES) ports cannot be SPAN or RSPAN source ports.

You can use the SPAN or RSPAN 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.


This section includes these topics:

- [Local SPAN, page 21-2](#)
- [Remote SPAN, page 21-2](#)
- [SPAN and RSPAN Concepts and Terminology, page 21-3](#)
- [SPAN and RSPAN Interaction with Other Features, page 21-8](#)

Local SPAN

Local SPAN supports a SPAN session entirely within one switch; all source ports or source VLANs and destination ports reside in the same switch. Local SPAN copies traffic from one or more source ports in any VLAN or from one or more VLANs to a destination port for analysis. For example, in [Figure 21-1](#), all traffic on port 5 (the source port) is mirrored to port 10 (the destination port). A network analyzer on port 10 receives all network traffic from port 5 without being physically attached to port 5.

Figure 21-1 Example of Local SPAN Configuration on a Switch




43580

Remote SPAN

RSPAN supports source ports, source VLANs, and destination ports on different switches, 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 RSPAN traffic from the source ports or VLANs is copied into the RSPAN VLAN and forwarded over trunk ports carrying the RSPAN VLAN to a destination session monitoring the RSPAN VLAN. Each RSPAN source switch must have either ports or VLANs as RSPAN sources. The destination is always a physical port.

Figure 21-2 shows RSPAN source ports configured on Switch A and Switch B with RSPAN destination ports configured on Switch C.

Figure 21-2 Example of RSPAN Configuration



SPAN and RSPAN Concepts and Terminology

This section describes concepts and terminology associated with SPAN and RSPAN configuration.

SPAN Sessions

SPAN sessions (local or remote) allow you to monitor traffic on one or more ports, or one or more VLANs, and send the monitored traffic to one or more destination ports.

A local SPAN session is an association of a destination port with source ports or source VLANs, all on a single network device. Local SPAN does not have separate source and destination sessions. Local SPAN sessions gather a set of ingress and egress packets specified by the user and form them into a stream of SPAN data, which is directed to the destination port.

RSPAN consists of at least one RSPAN source session, an RSPAN VLAN, and at least one RSPAN destination session. You separately configure RSPAN source sessions and RSPAN destination sessions on different network devices. To configure an RSPAN source session on a device, you associate a set of source ports or source VLANs with an RSPAN VLAN. The output of this session is the stream of SPAN

packets that are sent to the RSPAN VLAN. To configure an RSPAN destination session on another device, you associate the destination port with the RSPAN VLAN. The destination session collects all RSPAN VLAN traffic and sends it out the RSPAN destination port.

An RSPAN source session is very similar to a local SPAN session, except for where the packet stream is directed. In an RSPAN source session, SPAN packets are relabeled with the RSPAN VLAN ID and directed over normal trunk ports to the destination switch.

An RSPAN destination session takes all packets received on the RSPAN VLAN, strips off the VLAN tagging, and presents them on the destination port. Its purpose is to present a copy of all RSPAN VLAN packets (except Layer 2 control packets) to the user for analysis.

There can be more than one source session and more than one destination session active in the same RSPAN VLAN. There can also be intermediate switches separating the RSPAN source and destination sessions. These switches need not be capable of running RSPAN, but they must handle the requirements of the RSPAN VLAN (see the [“RSPAN VLAN” section on page 21-8](#)).

Traffic monitoring in a SPAN session has these restrictions:

- Sources can be ports or VLANs, but you cannot mix source ports and source VLANs in the same session.
- The switch supports up to two source sessions; you can run both a local SPAN and an RSPAN source session in the same switch. The switch supports a total of 66 source and RSPAN destination sessions.
- ES ports do not support SPAN source sessions.
- You can have multiple destination ports in a SPAN session, but no more than 64 destination ports.
- You can configure two separate SPAN or RSPAN source sessions with separate or overlapping sets of SPAN source ports and VLANs. Both switched and routed ports can be configured as SPAN sources and destinations.
- 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, can result in dropped or lost packets.
- When RSPAN is enabled, each packet being monitored is transmitted twice, once as normal traffic and once as a monitored packet. Therefore monitoring a large number of ports or VLANs could potentially generate large amounts of network traffic.
- 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 or VLAN for that session.
- The switch does not support a combination of local SPAN and RSPAN in a single session. That is, an RSPAN source session cannot have a local destination port, an RSPAN destination session cannot have a local source port, and an RSPAN destination session and an RSPAN source session that are using the same RSPAN VLAN cannot run on the same switch.

Monitored Traffic

SPAN sessions can monitor 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 or VLAN before any modification or processing is performed by the switch. A copy of each packet received by the source is sent to the destination port for that SPAN session.

Packets that are modified because of routing or quality of service (QoS)—for example, modified Differentiated Services Code Point (DSCP)—are copied before modification.

Features that can cause a packet to be dropped during receive processing have no effect on ingress SPAN; the destination port receives a copy of the packet even if the actual incoming packet is dropped. These features include IP standard and extended input access control lists (ACLs), ingress QoS policing, VLAN ACLs and egress QoS policing.

- **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.

Packets that are modified because of routing—for example, with modified time-to-live (TTL), MAC-address, or QoS values—are duplicated (with the modifications) at the destination port.

Features that can cause a packet to be dropped during transmit processing also affect the duplicated copy for SPAN. These features include IP standard and extended output ACLs and egress QoS policing.

- **Both**—In a SPAN session, you can also monitor a port or VLAN for both received and sent packets. This is the default.

The default configuration for local SPAN session ports is to send all packets untagged. SPAN also does not normally monitor bridge protocol data unit (BPDU) packets and Layer 2 protocols, such as Cisco Discovery Protocol (CDP), VLAN Trunk Protocol (VTP), Dynamic Trunking Protocol (DTP), Spanning Tree Protocol (STP), and Port Aggregation Protocol (PAgP). However, when you enter the **encapsulation replicate** keywords when configuring a destination port, these changes occur:

- Packets are sent on the destination port with the same encapsulation—untagged, IEEE 802.1Q, or Inter-Switch Link (ISL)—that they had on the source port.
- Packets of all types, including BPDU and Layer 2 protocol packets are monitored.

Therefore, a local SPAN session with encapsulation replicate enabled can have a mixture of untagged, 802.1Q, and ISL tagged packets appear on the destination port.

Switch congestion can cause packets to be dropped at ingress source ports, egress source ports, or SPAN destination ports. In general, these characteristics are independent of one another. For example:

- A packet might be forwarded normally but dropped from monitoring due to an oversubscribed SPAN destination port.
- An ingress packet might be dropped from normal forwarding, but still appear on the SPAN destination port.
- An egress packet dropped because of switch congestion is also dropped from egress SPAN.

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 Rx monitor on port A and Tx monitor on port B. If a packet enters the switch through port A and is switched to port B, both incoming and outgoing packets are sent to the destination port. Both packets are the same (unless a Layer-3 rewrite occurs, in which case the packets are different because of the packet modification).

Source Ports

A source port (also called a *monitored port*) is a switched or routed port that you monitor for network traffic analysis. In a local SPAN session or RSPAN source session, you can monitor source ports or VLANs for traffic in one or both directions. The switch supports any number of source ports (up to the maximum number of available ports on the switch) and any number of source VLANs (up to the maximum number of VLANs supported). However, the switch supports a maximum of two sessions (local or RSPAN) with source ports or VLANs and you cannot mix ports and VLANs in a single session.

A source port has these characteristics:

- It can be monitored in multiple SPAN sessions.
- Each source port can be configured with a direction (ingress, egress, or both) to monitor.
- It can be any port type (for example, EtherChannel, Fast Ethernet, Gigabit Ethernet, and so forth).
- For EtherChannel sources, you can monitor traffic for the entire EtherChannel or individually on a physical port as it participates in the port channel.
- It can be an access port, trunk port, routed port, or voice VLAN port.
- It cannot be an ES port.
- It cannot be a destination port.
- Source ports can be in the same or different VLANs.
- You can monitor multiple source ports in a single session.

Source VLANs

VLAN-based SPAN (VSPAN) is the monitoring of the network traffic in one or more VLANs. The SPAN or RSPAN source interface in VSPAN is a VLAN ID and traffic is monitored on all the ports for that VLAN.

VSPAN has these characteristics:

- All active ports in the source VLAN (except ES ports) are included as source ports and can be monitored in either or both directions.
- On a given port, only traffic on the monitored VLAN is sent to the destination port.
- If a destination port belongs to a source VLAN, it is excluded from the source list and is not monitored.
- If ports are added to or removed from the source VLANs, the traffic on the source VLAN received by those ports is added to or removed from the sources being monitored.
- You cannot use filter VLANs in the same session with VLAN sources.
- You can monitor only Ethernet VLANs.

VLAN Filtering

When you monitor a trunk port as a source port, by default, all VLANs active on the trunk are monitored. You can limit SPAN traffic monitoring on trunk source ports to specific VLANs by using VLAN filtering.

- VLAN filtering applies only to trunk ports or to voice VLAN ports.
- VLAN filtering applies only to port-based sessions and is not allowed in sessions with VLAN sources.
- When a VLAN filter list is specified, only those VLANs in the list are monitored on trunk ports or on voice VLAN access ports.
- SPAN traffic coming from other port types is not affected by VLAN filtering; that is, all VLANs are allowed on other ports.
- VLAN filtering affects only traffic forwarded to the destination SPAN port and does not affect the switching of normal traffic.

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 ports or VLANs and sends the SPAN packets to the user, usually a network analyzer.

A destination port has these characteristics:

- For a local SPAN session, the destination port must reside on the same switch as the source port. For an RSPAN session, it is located on the switch containing the RSPAN destination session. There is no destination port on a switch running only an RSPAN source session.
- When a port is configured as a SPAN destination port, the configuration overwrites the original port configuration. When the SPAN destination configuration is removed, the port reverts to its previous configuration. If a configuration change is made to the port while it is acting as a SPAN destination port, the change does not take effect until the SPAN destination configuration had been removed.
- If the port was in an EtherChannel group, it is removed from the group while it is a destination port. If it was a routed port, it is no longer a routed port.
- It can be any Ethernet physical port.
- It cannot be a secure port.
- It cannot be a source port.
- It cannot be an EtherChannel group or a VLAN.
- It can participate in only one SPAN session at a time (a destination port in one SPAN session cannot be a destination port for a second SPAN session).
- When it is active, incoming traffic is disabled. The port does not transmit any traffic except that required for the SPAN session. Incoming traffic is never learned or forwarded on a destination port.
- If ingress traffic forwarding is enabled for a network security device, the destination port forwards traffic at Layer 2.
- It does not participate in any of the Layer 2 protocols (STP, VTP, CDP, DTP, PagP).
- A destination port that belongs to a source VLAN of any SPAN session is excluded from the source list and is not monitored.
- The maximum number of destination ports in a switch is 64.

Local SPAN and RSPAN destination ports behave differently regarding VLAN tagging and encapsulation:

- For local SPAN, if the **encapsulation replicate** keywords are specified for the destination port, these packets appear with the original encapsulation (untagged, ISL, or 802.1Q). If these keywords are not specified, packets appear in the untagged format. Therefore, the output of a local SPAN session with **encapsulation replicate** enabled can contain a mixture of untagged, 802.1Q, or ISL tagged packets.
- For RSPAN, the original VLAN ID is lost because it is overwritten by the RSPAN VLAN identification. Therefore, all packets appear on the destination port as untagged.

RSPAN VLAN

The RSPAN VLAN carries SPAN traffic between RSPAN source and destination sessions. It has these special characteristics:

- All traffic in the RSPAN VLAN is always flooded.
- No MAC address learning occurs on the RSPAN VLAN.
- RSPAN VLAN traffic only flows on trunk ports.
- RSPAN VLANs must be configured in VLAN configuration mode by using the **remote-span** VLAN configuration mode command.
- STP can run on RSPAN VLAN trunks but not on SPAN destination ports.

For VLANs 1 to 1005 that are visible to VLAN Trunking Protocol (VTP), the VLAN ID and its associated RSPAN characteristic are propagated by VTP. If you assign an RSPAN VLAN ID in the extended VLAN range (1006 to 4094), you must manually configure all intermediate switches.

It is normal to have multiple RSPAN VLANs in a network at the same time with each RSPAN VLAN defining a network-wide RSPAN session. That is, multiple RSPAN source sessions anywhere in the network can contribute packets to the RSPAN session. It is also possible to have multiple RSPAN destination sessions throughout the network, monitoring the same RSPAN VLAN and presenting traffic to the user. The RSPAN VLAN ID separates the sessions.

SPAN and RSPAN Interaction with Other Features

SPAN interacts with these features:

- Routing—SPAN does not monitor routed traffic. VSPAN only monitors traffic that enters or exits the switch, not traffic that is routed between VLANs. For example, if a VLAN is being Rx-monitored and the switch routes traffic from another VLAN to the monitored VLAN, that traffic is not monitored and not received on the SPAN destination port.
- Spanning Tree Protocol (STP)—A destination port does not participate in STP while its SPAN or RSPAN session is active. The destination 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 or destination ports at any time. However, changes in VLAN membership or trunk settings for a destination port do not take effect until you remove the SPAN destination configuration. Changes in VLAN membership or trunk settings for a source port immediately take effect, and the respective SPAN sessions automatically adjust 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 physical 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, because there are no longer any ports in the group, there is no data to monitor.

A physical port that belongs to an EtherChannel group can be configured as a SPAN source port and still be a part of the EtherChannel. In this case, data from the physical port is monitored as it participates in the EtherChannel. However, if a physical port that belongs to an EtherChannel group is configured as a SPAN destination, 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 *inactive* or *standalone* state.

If a physical port that belongs to an EtherChannel group is a destination port and the EtherChannel group is a source, the port is removed from the EtherChannel group and from the list of monitored ports.

- 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.
- A secure port cannot be a SPAN destination port.

For SPAN sessions, do not enable port security on ports with monitored egress when ingress forwarding is enabled on the destination port. For RSPAN source sessions, do not enable port security on any ports with monitored egress.

- An 802.1x port can be a SPAN source port. You can enable 802.1X on a port that is a SPAN destination port; however, 802.1x is disabled until the port is removed as a SPAN destination.

For SPAN sessions, do not enable 802.1x on ports with monitored egress when ingress forwarding is enabled on the destination port. For RSPAN source sessions, do not enable 802.1x on any ports that are egress monitored.

Configuring SPAN and RSPAN

This section describes how to configure SPAN on your switch. It contains this configuration information:

- [Default SPAN and RSPAN Configuration, page 21-9](#)
- [Configuring Local SPAN, page 21-10](#)
- [Configuring RSPAN, page 21-16](#)

Default SPAN and RSPAN Configuration

[Table 21-1](#) shows the default SPAN and RSPAN configuration.

Table 21-1 Default SPAN and RSPAN Configuration

Feature	Default Setting
SPAN state (SPAN and RSPAN)	Disabled.
Source port traffic to monitor	Both received and sent traffic (both).
Encapsulation type (destination port)	Native form (untagged packets).
Ingress forwarding (destination port)	Disabled
VLAN filtering	On a trunk interface used as a source port, all VLANs are monitored.
RSPAN VLANs	None configured.

Configuring Local SPAN

This section describes how to configure Local SPAN on your switch. It contains this configuration information:

- [SPAN Configuration Guidelines, page 21-10](#)
- [Creating a Local SPAN Session, page 21-11](#)
- [Creating a Local SPAN Session and Configuring Ingress Traffic, page 21-13](#)
- [Specifying VLANs to Filter, page 21-15](#)

SPAN Configuration Guidelines

Follow these guidelines when configuring SPAN:

- You can configure a total of two local SPAN sessions or RSPAN source sessions on each switch. You can have a total of 66 SPAN sessions (local, RSPAN source, and RSPAN destination) on a switch.
- For SPAN sources, you can monitor traffic for a single port or VLAN or a series or range of ports or VLANs for each session. You cannot mix source ports and source VLANs within a single SPAN session.
- The destination port cannot be a source port; a source port cannot be a destination port.
- You cannot have two SPAN sessions using the same destination port.
- 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.
- Entering SPAN configuration commands does not remove previously configured SPAN parameters. You must enter the **no monitor session** {*session_number* | **all** | **local** | **remote**} global configuration command to delete configured SPAN parameters.
- For local SPAN, outgoing packets through the SPAN destination port carry the original encapsulation headers—untagged, ISL, or IEEE 802.1Q— if the **encapsulation replicate** keywords are specified. If the keywords are not specified, the packets are sent in native form. For RSPAN destination ports, outgoing packets are not tagged.
- 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 or source VLAN are enabled.
- You can limit SPAN traffic to specific VLANs by using the **filter vlan** keyword. If a trunk port is being monitored, only traffic on the VLANs specified with this keyword is monitored. By default, all VLANs are monitored on a trunk port.
- You cannot mix source VLANs and filter VLANs within a single SPAN session.

Creating a Local SPAN Session

Beginning in privileged EXEC mode, follow these steps to create a SPAN session and specify the source (monitored) ports or VLANs and the destination (monitoring) ports:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no monitor session { <i>session_number</i> all local remote }	Remove any existing SPAN configuration for the session. For <i>session_number</i> , the range is from 1 to 66. Specify all to remove all SPAN sessions, local to remove all local sessions, or remote to remove all remote SPAN sessions.
Step 3	monitor session <i>session_number</i> source { interface <i>interface-id</i> vlan <i>vlan-id</i> } [, -] [both rx tx]	Specify the SPAN session and the source port (monitored port). For <i>session_number</i> , the range is from 1 to 66. For <i>interface-id</i> , specify the source port or source VLAN to monitor. - For source <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>). Valid port channel numbers are 1 to 12. - For <i>vlan-id</i>, specify the source VLAN to monitor. The range is 1 to 4094 (excluding the RSPAN VLAN). Note A single session can include multiple sources (ports or VLANs), defined in a series of commands, but you cannot combine source ports and source VLANs in one session. (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 SPAN monitors both sent and received traffic. both—Monitor both received and sent traffic. This is the default. rx—Monitor received traffic. tx—Monitor sent traffic. Note You can use the monitor session <i>session_number</i> source command multiple times to configure multiple source ports.

	Command	Purpose
Step 4	monitor session <i>session_number</i> destination { interface <i>interface-id</i> [, -] [encapsulation replicate]}	Specify the SPAN session and the destination port (monitoring port). For <i>session_number</i> , specify the session number entered in step 3. Note For local SPAN, you must use the same session number for the source and destination interfaces. For <i>interface-id</i> , specify the destination port. The destination interface must be a physical port; it cannot be an EtherChannel, and it cannot be a VLAN. (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) Enter encapsulation replicate to specify that the destination interface replicates the source interface encapsulation method. If not selected, the default is to send packets in native form (untagged). Note You can use monitor session <i>session_number</i> destination command multiple times to configure multiple destination ports.
Step 5	end	Return to privileged EXEC mode.
Step 6	show monitor [session <i>session_number</i>] show running-config	Verify the configuration.
Step 7	copy running-config startup-config	(Optional) Save the configuration in the configuration file.

To delete a SPAN session, use the **no monitor session** *session_number* global configuration command. To remove a source or destination port or VLAN from the SPAN session, use the **no monitor session** *session_number* **source** {**interface** *interface-id* | **vlan** *vlan-id*} global configuration command or the **no monitor session** *session_number* **destination interface** *interface-id* global configuration command. For destination interfaces, the **encapsulation replicate** keywords are ignored with the **no** form of the command.

This example shows how to set up local SPAN session 1, for monitoring source port traffic to a destination port. First, any existing SPAN configuration for session 1 is deleted, and then bidirectional traffic is mirrored from source Gigabit Ethernet port 1 to destination Gigabit Ethernet port 2, retaining the encapsulation method.

```
Switch(config)# no monitor session 1
Switch(config)# monitor session 1 source interface gigabitethernet1/0/1
Switch(config)# monitor session 1 destination interface gigabitethernet1/0/2
encapsulation replicate
Switch(config)# end
```

This example shows how to remove a port as a SPAN source for SPAN session 1:

```
Switch(config)# no monitor session 1 source interface gigabitethernet1/0/1
Switch(config)# end
```

This example shows how to disable monitoring received traffic on a port that was configured for bidirectional monitoring:

```
Switch(config)# no monitor session 1 source interface gigabitethernet1/0/1 rx
```

Traffic received on port 1 is not monitored, but traffic sent from this port continues to be monitored.

This example shows how to remove any existing configuration on SPAN session 2, configure SPAN session 2 to monitor received traffic on all ports belonging to VLANs 1 through 3, and send it to a destination interface. The configuration is then modified to also monitor all traffic on all ports belonging to VLAN 10.

```
Switch(config)# no monitor session 2
Switch(config)# monitor session 2 source vlan 1 - 3 rx
Switch(config)# monitor session 2 destination interface gigabitethernet1/0/2
Switch(config)# monitor session 2 source vlan 10
Switch(config)# end
```

Creating a Local SPAN Session and Configuring Ingress Traffic

Beginning in privileged EXEC mode, follow these steps to create a SPAN session, to specify the source ports or VLANs and the destination ports, and to enable ingress traffic on the destination port for a network security device (such as a Cisco IDS Sensor Appliance).



Note

Refer to the [“Creating a Local SPAN Session” section on page 21-11](#) for details about the keywords not related to ingress traffic.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no monitor session { <i>session_number</i> all local remote }	Remove any existing SPAN configuration for the session.
Step 3	monitor session <i>session_number</i> source { interface <i>interface-id</i> vlan <i>vlan-id</i> } [, -] [both rx tx]	Specify the SPAN session and the source port (monitored port).

	Command	Purpose
Step 4	monitor session <i>session_number</i> destination { interface <i>interface-id</i> [, -] [encapsulation replicate] [ingress { dot1q vlan <i>vlan-id</i> isl untagged vlan <i>vlan-id</i> vlan <i>vlan-id</i> }}}	Specify the SPAN session, the destination port, the packet encapsulation, and the ingress VLAN and encapsulation. For <i>session_number</i> , specify the session number entered in step 3. For <i>interface-id</i> , specify the destination port. The destination interface must be a physical port; it cannot be an EtherChannel, and it cannot be a VLAN. (Optional) [, -] Specify a series or range of interfaces. Enter a space before and after the comma or hyphen. (Optional) Enter encapsulation replicate to specify that the destination interface replicates the source interface encapsulation method. If not selected, the default is to send packets in native form (untagged). Enter ingress with keywords to enable ingress traffic forwarding on the destination port and specify the encapsulation type: • dot1q vlan <i>vlan-id</i>—Forward ingress packets with 802.1Q encapsulation with the specified VLAN as the default VLAN. • isl—Forward ingress packets with ISL encapsulation. • untagged vlan <i>vlan-id</i> or vlan <i>vlan-id</i>—Forward ingress packets with untagged encapsulation type with the specified VLAN as the default VLAN.
Step 5	end	Return to privileged EXEC mode.
Step 6	show monitor [session <i>session_number</i>] show running-config	Verify the configuration.
Step 7	copy running-config startup-config	(Optional) Save the configuration in the configuration file.

To delete a SPAN session, use the **no monitor session** *session_number* global configuration command. To remove a source or destination port or VLAN from the SPAN session, use the **no monitor session** *session_number* **source** {**interface** *interface-id* | **vlan** *vlan-id*} global configuration command or the **no monitor session** *session_number* **destination interface** *interface-id* global configuration command. For destination interfaces, the encapsulation and ingress options are ignored with the **no** form of the command.

This example shows how to remove any existing configuration on SPAN session 2, configure SPAN session 2 to monitor received traffic on Gigabit Ethernet source port 1, and send it to destination Gigabit Ethernet port 2 with the same egress encapsulation type as the source port, and enable ingress forwarding with 802.1Q encapsulation and VLAN 6 as the default ingress VLAN.

```
Switch(config)# no monitor session 2
Switch(config)# monitor session 2 source gigabitethernet1/0/1 rx
Switch(config)# monitor session 2 destination interface gigabitethernet1/0/2 encapsulation
replicate ingress dot1q vlan 6
Switch(config)# end
```

Specifying VLANs to Filter

Beginning in privileged EXEC mode, follow these steps to limit SPAN source traffic to specific VLANs:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no monitor session { <i>session_number</i> all local remote }	Remove any existing SPAN configuration for the session. For <i>session_number</i> , the range is from 1 to 66. Specify all to remove all SPAN sessions, local to remove all local sessions, or remote to remove all remote SPAN sessions.
Step 3	monitor session <i>session_number</i> source interface <i>interface-id</i>	Specify the characteristics of the source port (monitored port) and SPAN session. For <i>session_number</i> , the range is from 1 to 66. For <i>interface-id</i> , specify the source port to monitor. The interface specified must already be configured as a trunk port.
Step 4	monitor session <i>session_number</i> filter vlan <i>vlan-id</i> [, -]	Limit the SPAN source traffic to specific VLANs. For <i>session_number</i> , enter the session number specified in Step 3. For <i>vlan-id</i> , the range is 1 to 4094. (Optional) Use a comma (,) to specify a series of VLANs, or use a hyphen (-) to specify a range of VLANs. Enter a space before and after the comma; enter a space before and after the hyphen.
Step 5	monitor session <i>session_number</i> destination { interface <i>interface-id</i> [, -] [encapsulation replicate]}	Specify the SPAN session and the destination port (monitoring port). For <i>session_number</i> , specify the session number entered in step 3. For <i>interface-id</i> , specify the destination port. The destination interface must be a physical port; it cannot be an EtherChannel, and it cannot be a VLAN. (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) Enter encapsulation replicate to specify that the destination interface replicates the source interface encapsulation method. If not selected, the default is to send packets in native form (untagged).
Step 6	end	Return to privileged EXEC mode.
Step 7	show monitor [session <i>session_number</i>] show running-config	Verify the configuration.
Step 8	copy running-config startup-config	(Optional) Save the configuration in the configuration file.

To monitor all VLANs on the trunk port, use the **no monitor session** *session_number* **filter** global configuration command.

This example shows how to remove any existing configuration on SPAN session 2, configure SPAN session 2 to monitor traffic received on Gigabit Ethernet trunk port 2, and send traffic for only VLANs 1 through 5 and 9 to destination Gigabit Ethernet port 1.

```
Switch(config)# no monitor session 2
Switch(config)# monitor session 2 source interface gigabitethernet1/0/2 rx
Switch(config)# monitor session 2 filter vlan 1 - 5 , 9
Switch(config)# monitor session 2 destination interface gigabitethernet1/0/1
Switch(config)# end
```

Configuring RSPAN

This section describes how to configure RSPAN on your switch. It contains this configuration information:

- [RSPAN Configuration Guidelines, page 21-16](#)
- [Configuring a VLAN as an RSPAN VLAN, page 21-17](#)
- [Creating an RSPAN Source Session, page 21-18](#)
- [Creating an RSPAN Destination Session, page 21-19](#)
- [Creating an RSPAN Destination Session, page 21-19](#)
- [Specifying VLANs to Filter, page 21-22](#)

RSPAN Configuration Guidelines

Follow these guidelines when configuring RSPAN:

- All the items in the [“SPAN Configuration Guidelines” section on page 21-10](#) apply to RSPAN.
- 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.
- You can apply an output access control list (ACL) to RSPAN traffic to selectively filter or monitor specific packets. Specify these ACLs on the RSPAN VLAN in the RSPAN source switches.
- For RSPAN configuration, you can distribute the source ports and the destination ports across multiple switches in your network.
- RSPAN does not support BPDU packet monitoring or other Layer 2 switch protocols.
- The RSPAN VLAN is configured only on trunk ports and not on access ports. To avoid unwanted traffic in RSPAN VLANs, make sure that the VLAN remote-span feature is supported in all the participating switches.
- Access ports (including voice VLAN ports) on the RSPAN VLAN are put in the inactive state.
- RSPAN VLANs are included as sources for port-based RSPAN sessions when source trunk ports have active RSPAN VLANs. RSPAN VLANs can also be sources in SPAN sessions. However, since the switch does not monitor spanned traffic, it does not support egress spanning of packets on any RSPAN VLAN identified as the destination of an RSPAN source session on the switch.
- You can configure any VLAN as an RSPAN VLAN as long as these conditions are met:
 - The same RSPAN VLAN is used for an RSPAN session in all the switches.
 - All participating switches support RSPAN.

- We recommend that you configure an RSPAN VLAN before you configure an RSPAN source or a 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.

Configuring a VLAN as an RSPAN VLAN

First create a new VLAN to be the RSPAN VLAN for the RSPAN session. You must create the RSPAN VLAN in all switches that will participate in RSPAN. If the RSPAN VLAN-ID is in the normal range (lower than 1005) and VTP is 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 extended-range VLANs (greater than 1005), you must configure RSPAN VLAN on both source and destination switches and any intermediate switches.

Use VTP pruning to get an efficient flow of RSPAN traffic, or manually delete the RSPAN VLAN from all trunks that do not need to carry the RSPAN traffic.

Beginning in privileged EXEC mode, follow these steps to create an RSPAN VLAN:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vlan <i>vlan-id</i>	Enter a VLAN ID to create a VLAN, or enter the VLAN ID of an existing VLAN, and enter VLAN configuration mode. The range is from 2 to 1001 and from 1006 to 4094. Note The RSPAN VLAN cannot be VLAN 1 (the default VLAN) or VLAN IDs 1002 through 1005 (reserved for Token Ring and FDDI VLANs).
Step 3	remote-span	Configure the VLAN as an RSPAN VLAN.
Step 4	end	Return to privileged EXEC mode.
Step 5	copy running-config startup-config	(Optional) Save the configuration in the configuration file.

To remove the remote SPAN characteristic from a VLAN and convert it back to a normal VLAN, use the **no remote-span** VLAN configuration command.

This example shows how to create RSPAN VLAN 901.

```
Switch(config)# vlan 901
Switch(config-vlan)# remote span
Switch(config-vlan)# end
```

Creating an RSPAN Source Session

Beginning in privileged EXEC mode, follow these steps to start an RSPAN source session and to specify the monitored source and the destination RSPAN VLAN:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no monitor session { <i>session_number</i> all local remote }	Remove any existing RSPAN configuration for the session. For <i>session_number</i> , the range is from 1 to 66. Specify all to remove all RSPAN sessions, local to remove all local sessions, or remote to remove all remote SPAN sessions.
Step 3	monitor session <i>session_number</i> source { interface <i>interface-id</i> vlan <i>vlan-id</i> } [, -] [both rx tx]	Specify the RSPAN session and the source port (monitored port). For <i>session_number</i> , the range is from 1 to 66. Enter a source port or source VLAN for the RSPAN session: - 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>). Valid port channel numbers are 1 to 12. - For <i>vlan-id</i>, specify the source VLAN to monitor. The range is 1 to 4094 (excluding the RSPAN VLAN). Note A single session can include multiple sources (ports or VLANs), defined in a series of commands, but you cannot combine source ports and source VLANs in one session. (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. both—Monitor both received and sent traffic. rx—Monitor received traffic. tx—Monitor sent traffic.
Step 4	monitor session <i>session_number</i> destination remote vlan <i>vlan-id</i>	Specify the RSPAN session and the destination RSPAN VLAN. For <i>session_number</i> , enter the number defined in Step 3. For <i>vlan-id</i> , specify the source RSPAN VLAN to monitor.
Step 5	end	Return to privileged EXEC mode.
Step 6	show monitor [session <i>session_number</i>] show running-config	Verify the configuration.
Step 7	copy running-config startup-config	(Optional) Save the configuration in the configuration file.

To delete a SPAN session, use the **no monitor session** *session_number* global configuration command.

To remove a source port or VLAN from the SPAN session, use the **no monitor session** *session_number* **source** {**interface** *interface-id* | **vlan** *vlan-id*} global configuration command. To remove the RSPAN VLAN from the session, use the **no monitor session** *session_number* **destination remote vlan** *vlan-id*.

This example shows how to remove any existing RSPAN configuration for session 1, configure RSPAN session 1 to monitor multiple source interfaces, and configure the destination as RSPAN VLAN 901.

```
Switch(config)# no monitor session 1
Switch(config)# monitor session 1 source interface fastethernet1/0/10 tx
Switch(config)# monitor session 1 source interface gigabitethernet1/0/1 rx
Switch(config)# monitor session 1 source interface gigabitethernet1/0/2
Switch(config)# monitor session 1 source interface port-channel 12
Switch(config)# monitor session 1 destination remote vlan 901
Switch(config)# end
```

Creating an RSPAN Destination Session

You configure the RSPAN destination session on a different switch; that is, not the switch on which the source session was configured.

Beginning in privileged EXEC mode, follow these steps to define the RSPAN VLAN on that switch, to create an RSPAN destination session, and to specify the source RSPAN VLAN and the destination port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vlan <i>vlan-id</i>	Enter the VLAN ID of the RSPAN VLAN created from the source switch, and enter VLAN configuration mode. Note If both switches are participating in VTP and the RSPAN VLAN ID is from 2 to 1005, Steps 2 through 4 are not required because the RSPAN VLAN ID is propagated through the VTP network.
Step 3	remote-span	Identify the VLAN as the RSPAN VLAN.
Step 4	exit	Return to global configuration mode.
Step 5	no monitor session { <i>session_number</i> all local remote }	Remove any existing RSPAN configuration for the session. For <i>session_number</i> , the range is from 1 to 66. Specify all to remove all RSPAN sessions, local to remove all local sessions, or remote to remove all remote SPAN sessions.
Step 6	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> , the range is from 1 to 66. For <i>vlan-id</i> , specify the source RSPAN VLAN to monitor.

	Command	Purpose
Step 7	monitor session <i>session_number</i> destination interface <i>interface-id</i>	Specify the RSPAN session and the destination interface. For <i>session_number</i> , enter the number defined in Step 6. Note In an RSPAN destination session, you must use the same session number for the source RSPAN VLAN and the destination port. For <i>interface-id</i> , specify the destination interface. The destination interface must be a physical interface. Note Though visible in the command-line help string, encapsulation replicate is not supported for RSPAN. The original VLAN ID is overwritten by the RSPAN VLAN ID, and all packets appear on the destination port as untagged.
Step 8	end	Return to privileged EXEC mode.
Step 9	show monitor [session <i>session_number</i>] show running-config	Verify the configuration.
Step 10	copy running-config startup-config	(Optional) Save the configuration in the configuration file.

To delete a SPAN session, use the **no monitor session** *session_number* global configuration command. To remove a destination port from the SPAN session, use the **no monitor session** *session_number* **destination interface** *interface-id* global configuration command. To remove the RSPAN VLAN from the session, use the **no monitor session** *session_number* **source remote vlan** *vlan-id*.

This example shows how to configure VLAN 901 as the source remote VLAN and Gigabit Ethernet port 2 as the destination interface:

```
Switch(config)# monitor session 1 source remote vlan 901
Switch(config)# monitor session 1 destination interface gigabitethernet1/0/2
Switch(config)# end
```

Creating an RSPAN Destination Session and Configuring Ingress Traffic

Beginning in privileged EXEC mode, follow these steps to create an RSPAN destination session, to specify the source RSPAN VLAN and the destination port, and to enable ingress traffic on the destination port for a network security device (such as a Cisco IDS Sensor Appliance).



Note

Refer to the [“Creating an RSPAN Destination Session” section on page 21-19](#) for details about the keywords not related to ingress traffic. This procedure assumes the RSPAN VLAN has already been configured.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no monitor session { <i>session_number</i> all local remote }	Remove any existing SPAN configuration for the session.

	Command	Purpose
Step 3	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> , the range is from 1 to 66. For <i>vlan-id</i> , specify the source RSPAN VLAN to monitor.
Step 4	monitor session <i>session_number</i> destination { interface <i>interface-id</i> [, -] [ingress { dot1q vlan <i>vlan-id</i> isl untagged vlan <i>vlan-id</i> vlan <i>vlan-id</i> }]}	Specify the SPAN session, the destination port, the packet encapsulation, and the ingress VLAN and encapsulation. For <i>session_number</i> , enter the number defined in Step 4. Note In an RSPAN destination session, you must use the same session number for the source RSPAN VLAN and the destination port. For <i>interface-id</i> , specify the destination interface. The destination interface must be a physical interface. Note Though visible in the command-line help string, encapsulation replicate is not supported for RSPAN. The original VLAN ID is overwritten by the RSPAN VLAN ID, and all packets appear on the destination port as untagged. (Optional) [, -] Specify a series or range of interfaces. Enter a space before and after the comma; enter a space before and after the hyphen. Enter ingress with additional keywords to enable ingress traffic forwarding on the destination port and to specify the encapsulation type: • dot1q vlan <i>vlan-id</i>—Forward ingress packets with 802.1Q encapsulation with the specified VLAN as the default VLAN. • isl—Forward ingress packets with ISL encapsulation. • untagged vlan <i>vlan-id</i> or vlan <i>vlan-id</i>—Forward ingress packets with untagged encapsulation type with the specified VLAN as the default VLAN.
Step 5	end	Return to privileged EXEC mode.
Step 6	show monitor [session <i>session_number</i>] show running-config	Verify the configuration.
Step 7	copy running-config startup-config	(Optional) Save the configuration in the configuration file.

To delete an RSPAN session, use the **no monitor session** *session_number* global configuration command. To remove a destination port from the RSPAN session, use the **no monitor session** *session_number* **destination interface** *interface-id* global configuration command. The ingress options are ignored with the **no** form of the command.

This example shows how to configure VLAN 901 as the source remote VLAN in RSPAN session 2, to configure Gigabit Ethernet source port 2 as the destination interface, and to enable ingress forwarding on the interface with VLAN 6 as the default ingress VLAN.

```
Switch(config)# monitor session 2 source remote vlan 901
Switch(config)# monitor session 2 destination interface gigabitethernet1/0/2 ingress
vlan 6
Switch(config)# end
```

Specifying VLANs to Filter

Beginning in privileged EXEC mode, follow these steps to configure the RSPAN source session to limit RSPAN source traffic to specific VLANs:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no monitor session { <i>session_number</i> all local remote }	Remove any existing SPAN configuration for the session. For <i>session_number</i> , the range is from 1 to 66. Specify all to remove all SPAN sessions, local to remove all local sessions, or remote to remove all remote SPAN sessions.
Step 3	monitor session <i>session_number</i> source interface <i>interface-id</i>	Specify the characteristics of the source port (monitored port) and SPAN session. For <i>session_number</i> , the range is from 1 to 66. For <i>interface-id</i> , specify the source port to monitor. The interface specified must already be configured as a trunk port.
Step 4	monitor session <i>session_number</i> filter vlan <i>vlan-id</i> [, -]	Limit the SPAN source traffic to specific VLANs. For <i>session_number</i> , enter the session number specified in step 3. For <i>vlan-id</i> , the range is 1 to 4094. (Optional) Use a comma (,) to specify a series of VLANs or use a hyphen (-) to specify a range of VLANs. Enter a space before and after the comma; enter a space before and after the hyphen.
Step 5	monitor session <i>session_number</i> destination remote vlan <i>vlan-id</i>	Specify the RSPAN session and the destination remote VLAN (RSPAN VLAN). For <i>session_number</i> , enter the session number specified in step 3. For <i>vlan-id</i> , specify the RSPAN VLAN to carry the monitored traffic to the destination port.
Step 6	end	Return to privileged EXEC mode.
Step 7	show monitor [session <i>session_number</i>] show running-config	Verify the configuration.
Step 8	copy running-config startup-config	(Optional) Save the configuration in the configuration file.

To monitor all VLANs on the trunk port, use the **no monitor session** *session_number* **filter vlan** global configuration command.

This example shows how to remove any existing configuration on RSPAN session 2, configure RSPAN session 2 to monitor traffic received on a trunk port, and send traffic for only VLANs 1 through 5 and 9 to destination RSPAN VLAN 902.

```
Switch(config)# no monitor session 2
Switch(config)# monitor session 2 source interface gigabitethernet1/0/2 rx
Switch(config)# monitor session 2 filter vlan 1 - 5 , 9
Switch(config)# monitor session 2 destination remote vlan 902
Switch(config)# end
```

Displaying SPAN and RSPAN Status

To display the current SPAN or RSPAN configuration, use the **show monitor** user EXEC command. You can also use the **show running-config** privileged EXEC command to display configured SPAN or RSPAN sessions.



Configuring RMON

This chapter describes how to configure Remote Network Monitoring (RMON) on the Catalyst 3750 Metro 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 IOS Configuration Fundamentals Command Reference for Release 12.1*.

This chapter consists of these sections:

- [Understanding RMON, page 22-1](#)
- [Configuring RMON, page 22-2](#)
- [Displaying RMON Status, page 22-6](#)

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 22-1 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 a port.
- History (RMON group 2)—Collects a history group of statistics on Ethernet, Fast Ethernet, and Gigabit Ethernet ports for a specified polling interval.
- Alarm (RMON group 3)—Monitors a specific 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)—Specifies 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 software release use hardware counters for RMON data processing, the monitoring is more efficient, and little processing power is required.

Configuring RMON

These sections describe how to configure RMON on your switch:

- [Default RMON Configuration, page 22-3](#)
- [Configuring RMON Alarms and Events, page 22-3](#) (required)
- [Collecting Group History Statistics on an Interface, page 22-5](#) (optional)
- [Collecting Group Ethernet Statistics on an Interface, page 22-6](#) (optional)

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 24, “Configuring SNMP.”](#)

Beginning in privileged EXEC mode, follow these steps to enable RMON alarms and events. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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>] [owner <i>string</i>]	Set an alarm on a MIB object. - 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 values 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 <i>string</i>, specify the owner of the alarm.

	Command	Purpose
Step 3	rmon event <i>number</i> [description <i>string</i>] [log] [owner <i>string</i>] [trap <i>community</i>]	Add an event in the RMON event table that is associated with an RMON event number. - For <i>number</i>, assign an event number. The range is 1 to 65535. (Optional) For description <i>string</i>, 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 <i>string</i>, specify the owner of this event. (Optional) For trap <i>community</i>, enter the SNMP community string used for this trap.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 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
```


Collecting Group History Statistics 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. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface on which to collect history, and enter interface configuration mode.
Step 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. - 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. The range is 1 to 3600. The default is 1800 seconds. (Optional) For owner <i>ownername</i>, enter the name of the owner of the RMON group of statistics.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	show rmon history	Display the contents of the switch history table.
Step 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.

Collecting Group Ethernet Statistics on an Interface

Beginning in privileged EXEC mode, follow these steps to collect group Ethernet statistics on an interface. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface on which to collect statistics, and enter interface configuration mode.
Step 3	rmon collection stats <i>index</i> [owner <i>ownername</i>]	Enable RMON statistic collection on the interface. - 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	show rmon statistics	Display the contents of the switch statistics table.
Step 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.

This example shows how to collect RMON statistics for the owner *root*:

```
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# rmon collection stats 2 owner root
```

Displaying RMON Status

To display the RMON status, use one or more of the privileged EXEC commands in [Table 22-1](#):

Table 22-1 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 IOS Configuration Fundamentals Command Reference for Release 12.1*.



Configuring System Message Logging

This chapter describes how to configure system message logging on the Catalyst 3750 Metro switch.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

This chapter consists of these sections:

- [Understanding System Message Logging, page 23-1](#)
- [Configuring System Message Logging, page 23-2](#)
- [Displaying the Logging Configuration, page 23-12](#)

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 console.



Note

The syslog format is compatible with 4.3 BSD UNIX.

When the logging process is disabled, messages are sent only to the console. 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 console 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 consoles and each of the destinations. You can time-stamp 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 on the switch. If the switch fails, the log is lost unless you had saved it to flash memory.

You can remotely monitor system messages by viewing the logs on a syslog server or by accessing the switch through Telnet or through the console port.

Configuring System Message Logging

These sections describe how to configure system message logging:

- [System Log Message Format, page 23-2](#)
- [Default System Message Logging Configuration, page 23-3](#)
- [Disabling Message Logging, page 23-4](#) (optional)
- [Setting the Message Display Destination Device, page 23-4](#) (optional)
- [Synchronizing Log Messages, page 23-5](#) (optional)
- [Enabling and Disabling Timestamps on Log Messages, page 23-7](#) (optional)
- [Enabling and Disabling Sequence Numbers in Log Messages, page 23-7](#) (optional)
- [Defining the Message Severity Level, page 23-8](#) (optional)
- [Limiting Syslog Messages Sent to the History Table and to SNMP, page 23-9](#) (optional)
- [Configuring UNIX Syslog Servers, page 23-10](#) (optional)

System Log Message Format

System log messages can contain up to 80 characters and a percent sign (%), which follows the optional sequence number or time-stamp 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 23-1](#) describes the elements of syslog messages.

Table 23-1 System Log Message Elements

Element	Description
<i>seq no:</i>	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 23-7 .
<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 23-7 .
<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 23-4 on page 23-12 .
<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 23-3 on page 23-9 .

Table 23-1 System Log Message Elements (continued)

Element	Description
<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 GigabitEthernet1/0/1, changed state to up
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet1/0/2, 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 GigabitEthernet1/0/1, 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 23-2](#) shows the default system message logging configuration.

Table 23-2 Default System Message Logging Configuration

Feature	Default Setting
System message logging to the console	Enabled.
Console severity	Debugging (and numerically lower levels; see Table 23-3 on page 23-9).
Logging file configuration	No filename specified.
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 23-4 on page 23-12).
Server severity	Informational (and numerically lower levels; see Table 23-3 on page 23-9).

Disabling Message Logging

Message logging is enabled by default. It must be enabled to send messages to any destination other than the console. 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. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no logging on	Disable message logging.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config or show logging	Verify your entries.
Step 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 console before continuing. When the logging process is disabled, messages are displayed on the console 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 console. When this command is enabled, messages appear only after you press Return. For more information, see the [“Synchronizing Log Messages”](#) section on page 23-5.

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 in addition to the console. Beginning in privileged EXEC mode, use one or more of the following commands to specify the locations that receive messages. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	logging buffered [size]	Log messages to an internal buffer on the switch. The default buffer size is 4096. The range is 4096 to 2147483647 bytes. If the switch, the log file is lost unless you previously saved it to flash memory. See Step 4. 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.

	Command	Purpose
Step 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 23-10.
Step 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 on the switch. - 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 4096 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 23-3 on page 23-9. By default, the log file receives debugging messages and numerically lower levels.
Step 5	end	Return to privileged EXEC mode.
Step 6	terminal monitor	Log messages to a nonconsole 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.
Step 7	show running-config	Verify your entries.
Step 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.

To disable logging to the console, use the **no logging console** global configuration command. To disable logging to a file, use the **no logging file** [*severity-level-number* | *type*] global configuration command.

Synchronizing Log Messages

You can synchronize unsolicited messages and **debug** privileged EXEC command output with solicited device output and prompts for a specific console 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 configure the maximum number of buffers for storing asynchronous messages for the terminal after which messages are dropped.

When synchronous logging of unsolicited messages and **debug** command output is enabled, unsolicited device output appears on the console or printed after solicited device output appears or printed.

Unsolicited messages and **debug** command output appears on the console 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 console again displays the user prompt.

Beginning in privileged EXEC mode, follow these steps to configure synchronous logging. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	line [console vty] <i>line-number</i> [<i>ending-line-number</i>]	Specify the line to be configured for synchronous logging of messages. - Use the console keyword for configurations that occur through the switch console 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. You can change the setting of all 16 vty lines at once by entering: line vty 0 15 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: line vty 2 When you enter this command, the mode changes to line configuration.
Step 3	logging synchronous [level <i>severity-level</i> all] limit <i>number-of-buffers</i>	Enable synchronous logging of messages. (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 range is 0 to 2147483647. The default is 20.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 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 time-stamped.

Beginning in privileged EXEC mode, follow these steps to enable time-stamping of log messages. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	service timestamps log uptime or service timestamps log datetime [msec] [localtime] [show-timezone]	Enable log time-stamps. The first command enables time-stamps on log messages, showing the time since the system was rebooted. The second command enables time-stamps 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable time-stamps 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. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	service sequence-numbers	Enable sequence numbers.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 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 23-3](#).

Beginning in privileged EXEC mode, follow these steps to define the message severity level. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	logging console <i>level</i>	Limit messages logged to the console. By default, the console receives debugging messages and numerically lower levels (see Table 23-3 on page 23-9).
Step 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 23-3 on page 23-9).
Step 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 23-3 on page 23-9). For complete syslog server configuration steps, see the “ Configuring UNIX Syslog Servers ” section on page 23-10.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config or show logging	Verify your entries.
Step 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 console, use the **no logging console** global configuration command. To disable logging to a terminal other than the console, use the **no logging monitor** global configuration command. To disable logging to syslog servers, use the **no logging trap** global configuration command.

[Table 23-3](#) describes the *level* keywords. It also lists the corresponding UNIX syslog definitions from the most severe level to the least severe level.

Table 23-3 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 the Technical Assistance Center.
- 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 23-3 on page 23-9](#)) 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. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	logging history level¹	Change the default level of syslog messages stored in the history file and sent to the SNMP server. See Table 23-3 on page 23-9 for a list of <i>level</i> keywords. By default, warnings, errors, critical, alerts, and emergencies messages are sent.

	Command	Purpose
Step 3	logging history size <i>number</i>	Specify the number of syslog messages that can be stored in the history table. The default is to store one message. The range is 0 to 500 messages.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

1. Table 23-3 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. This procedure is optional.

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 decide what options must be added to or removed from the syslog command line to enable logging of remote syslog messages.

Step 1 Add a line such as the following to the file `/etc/syslog.conf`:

```
local7.debug /usr/adm/logs/cisco.log
```

The **local7** keyword specifies the logging facility to be used; see [Table 23-4 on page 23-12](#) for information on the facilities. The **debug** keyword specifies the syslog level; see [Table 23-3 on page 23-9](#) 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.

Step 2 Create the log file by entering these commands at the UNIX shell prompt:

```
$ touch /var/log/cisco.log
$ chmod 666 /var/log/cisco.log
```

Step 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. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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.
Step 3	logging trap level	Limit messages logged to the syslog servers. By default, syslog servers receive informational messages and lower. See Table 23-3 on page 23-9 for <i>level</i> keywords.
Step 4	logging facility facility-type	Configure the syslog facility. See Table 23-4 on page 23-12 for <i>facility-type</i> keywords. The default is local7 .
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 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 23-4](#) lists the UNIX system facilities supported by the software. For more information about these facilities, consult the operator's manual for your UNIX operating system.

Table 23-4 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
sys9-14	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 IOS Configuration Fundamentals Command Reference for Release 12.1*.



Configuring SNMP

This chapter describes how to configure the Simple Network Management Protocol (SNMP) on the Catalyst 3750 Metro 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 IOS Configuration Fundamentals Command Reference for Release 12.1*.

This chapter consists of these sections:

- [Understanding SNMP, page 24-1](#)
- [Configuring SNMP, page 24-5](#)
- [Displaying SNMP Status, page 24-15](#)

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 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 Transmission Control Protocol (TCP) connection, loss of connection to a neighbor, or other significant events.

This section includes information about these topics:

- [SNMP Versions, page 24-2](#)
- [SNMP Manager Functions, page 24-3](#)
- [SNMP Agent Functions, page 24-4](#)
- [SNMP Community Strings, page 24-4](#)
- [Using SNMP to Access MIB Variables, page 24-4](#)
- [SNMP Notifications, page 24-5](#)

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—ensuring that a packet was not tampered with in transit
 - Authentication—determining that the message is from a valid source
 - Encryption—mixing 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 crypto (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 24-1 identifies the characteristics of the different combinations of security models and levels.

Table 24-1 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 using SNMPv1, and SNMPv2C, and SNMPv3 protocols.

SNMP Manager Functions

The SNMP manager uses information in the MIB to perform the operations described in Table 24-2.

Table 24-2 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 or module 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

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 24-1](#), 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 24-1 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.



Note

SNMPv1 does not support 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.

Configuring SNMP

This section describes how to configure SNMP on your switch. It contains this configuration information:

- [Default SNMP Configuration, page 24-6](#)
- [SNMP Configuration Guidelines, page 24-6](#)
- [Disabling the SNMP Agent, page 24-7](#)
- [Configuring Community Strings, page 24-7](#)
- [Configuring SNMP Groups and Users, page 24-8](#)
- [Configuring SNMP Notifications, page 24-10](#)
- [Setting the Agent Contact and Location Information, page 24-13](#)
- [Limiting TFTP Servers Used Through SNMP, page 24-13](#)
- [SNMP Examples, page 24-14](#)

Default SNMP Configuration

Table 24-3 shows the default SNMP configuration.

Table 24-3 Default SNMP Configuration

Feature	Default Setting
SNMP agent	Enabled
SNMP community strings	Read-Only: Public Read-Write: Private Read-Write-all: Secret
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 IOS Configuration Fundamentals Command Reference for Release 12.1* 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 the engine ID 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no snmp-server	Disable the SNMP agent operation.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 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 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	snmp-server community <i>string</i> [view <i>view-name</i>] [ro rw] [<i>access-list-number</i>]	Configure the community string. • 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.

	Command	Purpose
Step 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. - 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 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 (engine ID) 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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. - 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 has 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 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>]	Configure a new SNMP group on the remote device. - For <i>groupname</i>, specify the name of the group. - Specify a security model: 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: 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>). Note The priv keyword is available only when the crypto software image is installed. (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.

	Command	Purpose
Step 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. - 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, or v2c, or v3). If you enter v3, you have these additional options: 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.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 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 IOS 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 24-4 describes the supported switch traps (notification types). You can enable any or all of these traps and configure a trap manager to receive them.



Note

Although visible in the command-line interface (CLI) online help, the **fru-ctrl** keyword is not supported.

Table 24-4 Switch Notification Types

Notification Type Keyword	Description
bgp	Generates BGP state change traps. This option is only available when the enhanced multilayer image is installed.
bridge	Generates STP bridge MIB traps.
cluster	Generates a trap when the cluster configuration changes.
config	Generates a trap for SNMP configuration changes.
config-copy	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.
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.
port-security	Generates SNMP port security traps. You can also set a maximum trap rate per second. The range is from 0 to 1000; the default is 0, which means that there is no rate limit.
rtr	Generates a trap for the SNMP Response Time Reporter (RTR).
snmp	Generates a trap for SNMP-type notifications for authentication, cold start, warm start, link up or link down.
stpx	Generates SNMP STP Extended MIB traps.
syslog	Generates SNMP syslog traps.
tty	Generates a trap for TCP connections.
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.

**Note**

Though visible in the command-line help string, the **fru-ctrl** and flash **insertion** and **removal** keywords are not supported.

You can use the **snmp-server host** global configuration command to a specific host to receive the notification types listed in [Table 24-4](#).

Beginning in privileged EXEC mode, follow these steps to configure the switch to send traps or informs to a host:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	snmp-server engineID remote <i>ip-address engineid-string</i>	Specify the engine ID for the remote host.
Step 3	snmp-server user <i>username</i> <i>groupname remote host [udp-port</i> <i>port] {v1 v2c v3 [auth {md5 sha}</i> <i>auth-password]] [encrypted] [access</i> <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.
Step 4	snmp-server host <i>host-addr</i> [traps informs] [version {1 2c 3 [auth noauth priv]] <i>community-string [udp-port port]</i> <i>[notification-type]</i>	Specify the recipient of an SNMP trap operation. - 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 does not support informs. (Optional) For version 3, select authentication level auth, noauth, or priv. Note The priv keyword is available only when the crypto software image is installed. - For <i>community-string</i>, enter the password-like community string sent with the notification operation. (Optional) For udp-port port, enter the remote device UDP port. (Optional) For <i>notification-type</i>, use the keywords listed in Table 24-4 on page 24-11. If no type is specified, all notifications are sent.
Step 5	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 24-4 on page 24-11 , 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.
Step 6	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.
Step 7	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.
Step 8	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.
Step 9	end	Return to privileged EXEC mode.

	Command	Purpose
Step 10	show running-config	Verify your entries.
Step 11	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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	snmp-server contact <i>text</i>	Set the system contact string. For example: snmp-server contact Dial System Operator at beeper 21555.
Step 3	snmp-server location <i>text</i>	Set the system location string. For example: snmp-server location Building 3/Room 222
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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.

	Command	Purpose
Step 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. - 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 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 *cisco.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 cisco.com version 2c public
```

This example shows how to send Entity MIB traps to the host *cisco.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 *cisco.com*.

```
Switch(config)# snmp-server enable traps entity
Switch(config)# snmp-server host cisco.com restricted entity
```

This example shows how to enable the switch to send all traps to the host *myhost.cisco.com* using the community string *public*:

```
Switch(config)# snmp-server enable traps
Switch(config)# snmp-server host myhost.cisco.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 24-5](#) to display SNMP information. For information about the fields in the output displays, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

Table 24-5 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.



Note

Though visible in the command-line help strings, the **fru-ctrl** keyword is not supported. The **snmp-server enable informs** command is not supported. To enable the sending of SNMP inform notifications, use the **snmp-server enable traps** command combined with the **snmp-server host host-addr informs** command.



Configuring Network Security with ACLs

This chapter describes how to configure network security on the Catalyst 3750 Metro switch by using access control lists (ACLs), which are also referred to in commands and tables as access lists.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release and the “Configuring IP Services” section of the *Cisco IOS IP and IP Routing Configuration Guide* and the *Cisco IOS IP and IP Routing Command Reference for IOS Release 12.1*.

This chapter consists of these sections:

- [Understanding ACLs, page 25-1](#)
- [Configuring IP ACLs, page 25-6](#)
- [Creating Named MAC Extended ACLs, page 25-25](#)
- [Configuring VLAN Maps, page 25-27](#)
- [Using VLAN Maps with Router ACLs, page 25-34](#)
- [Displaying ACL Configuration, page 25-38](#)

Understanding ACLs

Packet filtering can help limit network traffic and restrict network use by certain users or devices. ACLs can filter traffic as it passes through a router or switch and permit or deny packets crossing specified interfaces or VLANs. 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. It tests packets against the conditions in an access list one by one. The first match determines whether the switch accepts or rejects the packets. 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 packets. If there are no restrictions, the switch forwards the packet; otherwise, the switch drops the packet. The switch can access-control all packets it switches, including packets bridged within a VLAN.

You configure access lists on a router or Layer 3 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 router interfaces. For example, you can allow e-mail traffic to be forwarded but not Telnet traffic. ACLs can be configured to block inbound traffic, outbound traffic, or both.

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 IP ACLs and Ethernet (MAC) ACLs:

- IP ACLs filter IP traffic, including TCP, User Datagram Protocol (UDP), Internet Group Management Protocol (IGMP), and Internet Control Message Protocol (ICMP).
- Ethernet ACLs filter non-IPv4 traffic.

This switch also supports quality of service (QoS) classification ACLs. For more information, see the [“Ingress Classification Based on QoS ACLs” section on page 26-8](#).

This section includes information on these topics:

- [Supported ACLs, page 25-2](#)
- [Handling Fragmented and Unfragmented Traffic, page 25-5](#)

Supported ACLs

The switch supports three applications of ACLs to filter traffic:

- Router ACLs access-control routed traffic between VLANs and are applied to Layer 3 interfaces.
- Port ACLs access-control traffic entering a Layer 2 interface. The switch does not support port ACLs in the outbound direction. You can apply only one IP access list and one MAC access list to a Layer 2 interface.
- VLAN ACLs or VLAN maps access-control all packets (bridged and routed). You can use VLAN maps to filter traffic between devices in the same VLAN. VLAN maps are configured to provide access-control based on Layer 3 addresses for IP. Unsupported protocols are access-controlled through MAC addresses using Ethernet ACEs. After a VLAN map is applied to a VLAN, all packets (routed or bridged) entering the VLAN are checked against the VLAN map. Packets can either enter the VLAN through a switch port or through a routed port after being routed.

You can use router ACLs, input port ACLs, and VLAN maps on the same switch. However, a port ACL takes precedence over a router ACL or VLAN map.

- When both an input port ACL and a VLAN map are applied, incoming packets received on ports with a port ACL applied are filtered by the port ACL. Other packets are filtered by the VLAN map.
- When an input router ACL and input port ACL exist in an switch virtual interface (SVI), incoming packets received on ports to which a port ACL is applied are filtered by the port ACL. Incoming routed IP packets received on other ports are filtered by the router ACL. Other packets are not filtered.
- When an output router ACL and input port ACL exist in an SVI, incoming packets received on the ports to which a port ACL is applied are filtered by the port ACL. Outgoing routed IP packets are filtered by the router ACL. Other packets are not filtered.
- When a VLAN map, input router ACL, and input port ACL exist in an SVI, incoming packets received on the ports to which a port ACL is applied are only filtered by the port ACL. Incoming routed IP packets received on other ports are filtered by both the VLAN map and the router ACL. Other packets are filtered only by the VLAN map.
- When a VLAN map, output router ACL, and input port ACL exist in an SVI, incoming packets received on the ports to which a port ACL is applied are only filtered by the port ACL. Outgoing routed IP packets are filtered by both the VLAN map and the router ACL. Other packets are filtered only by the VLAN map.

Router ACLs

You can apply router ACLs on SVIs, which are Layer 3 interfaces to VLANs; on physical Layer 3 interfaces; and on Layer 3 EtherChannel interfaces. You apply router ACLs on interfaces for specific directions (inbound or outbound). You can apply one router ACL in each direction on an interface.


One ACL can be used with multiple features for a given interface, and one feature can use multiple ACLs. When a single router ACL is used by multiple features, it is examined multiple times.

- 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 ACLs associated with features configured on a given interface and a direction. As packets enter the switch on an interface, ACLs associated with all inbound features configured on that interface are examined. After packets are routed and before they are forwarded to the next hop, all ACLs associated with outbound features configured on the egress interface are examined.

ACLs permit or deny packet forwarding based on how the packet matches the entries in the ACL, and can be used to control access to a network or to part of a network. In [Figure 25-1](#), 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 25-1 Using ACLs to Control Traffic to a Network



Port ACLs

Port ACLs are ACLs that are applied to Layer 2 interfaces on a switch. Port ACLs are supported only on physical interfaces and not on EtherChannel interfaces. Port ACLs are applied only on interfaces for inbound traffic.

These access lists are supported on Layer 2 interfaces:

- Standard IP access lists using source addresses
- Extended IP access lists using source and destination addresses and optional protocol type information
- MAC extended access lists using source and destination MAC addresses and optional protocol type information

As with router ACLs, the switch examines ACLs associated with features configured on a given interface and permits or denies packet forwarding based on how the packet matches the entries in the ACL. ACLs can only be applied to Layer 2 interfaces in the inbound direction. In the example in [Figure 25-1](#), if all workstations were in the same VLAN, ACLs applied at the Layer 2 input would allow Host A to access the Human Resources network, but prevent Host B from accessing the same network.

When you apply a port ACL to a trunk port, the ACL filters traffic on all VLANs present on the trunk port. When you apply a port ACL to a port with voice VLAN, the ACL filters traffic on both data and voice VLANs.

With port ACLs, you can filter IPv4 traffic by using IP access lists and non-IPv4 traffic by using MAC addresses. You can filter both IP and non-IP traffic on the same Layer 2 interface by applying both an IP access list and a MAC access list to the interface.



Note

You cannot apply more than one IP access list and one MAC access list to a Layer 2 interface. If an IP access list or MAC access list is already configured on a Layer 2 interface and you apply a new IP access list or MAC access list to the interface, the new ACL replaces the previously configured one.


VLAN Maps

VLAN ACLs or VLAN maps can access-control *all* traffic. You can apply VLAN maps to all packets that are routed into or out of a VLAN or are bridged within a VLAN. VLAN maps are used for security packet filtering. VLAN maps are not defined by direction (input or output).

You can configure VLAN maps to match Layer 3 addresses for IP traffic. All non-IPv4 protocols are access-controlled through MAC addresses and Ethertype using MAC VLAN maps. (IP traffic *is not* access controlled by MAC VLAN maps.) You can enforce VLAN maps only on packets going through the switch; you cannot enforce VLAN maps on traffic between hosts on a hub or on another switch connected to this switch.

With VLAN maps, forwarding of packets is permitted or denied, based on the action specified in the map. [Figure 25-2](#) illustrates how a VLAN map is applied to deny a specific type of traffic from Host A in VLAN 10 from being forwarded. You can apply only one VLAN map to a VLAN.

Figure 25-2 Using VLAN Maps to Control Traffic



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, 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 permit tcp any host 10.1.1.2
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. Instead, they match the third ACE (a permit).

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 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 fourth ACE (a deny). All other fragments also match the fourth 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.

Configuring IP ACLs

Configuring IP ACLs on the switch is the same as configuring IP ACLs on other Cisco switches and routers. The process is briefly described here. For more detailed information on configuring ACLs, refer to the “Configuring IP Services” chapter in the *Cisco IP and IP Routing Configuration Guide for IOS Release 12.1*. For detailed information about the commands, refer to *Cisco IOS IP and IP Routing Command Reference for IOS Release 12.1*.

The switch does not support these IOS router ACL-related features:

- Non-IPv4 protocol ACLs (see [Table 25-1 on page 25-7](#)) or bridge-group ACLs
- IP accounting
- Inbound and outbound rate limiting (except with QoS ACLs)
- Reflexive ACLs or dynamic ACLs (except for some specialized dynamic ACLs used by the switch clustering feature)
- ACL logging for port ACLs and VLAN maps

These are the steps to use IP ACLs on the switch:

-
Step 1	Create an ACL by specifying an access list number or name and access conditions.
Step 2	Apply the ACL to interfaces or terminal lines. You can also apply standard and extended IP ACLs to VLAN maps.
-

This section includes the following information:

- [Creating Standard and Extended IP ACLs, page 25-6](#)
- [Applying an IP ACL to a Terminal Line, page 25-17](#)
- [Applying an IP ACL to an Interface, page 25-18](#)
- [Hardware and Software Treatment of IP ACLs, page 25-19](#)
- [IP ACL Configuration Examples, page 25-20](#)

Creating Standard and Extended IP ACLs

This section describes IP ACLs. An ACL is a sequential collection of permit and deny conditions. 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.

The software supports these types of ACLs or access lists for IP:

- 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.

These sections describe access lists and how to create them:

- [Access List Numbers, page 25-7](#)
- [Creating a Numbered Standard ACL, page 25-8](#)
- [Creating a Numbered Extended ACL, page 25-9](#)
- [Creating Named Standard and Extended ACLs, page 25-13](#)
- [Using Time Ranges with ACLs, page 25-15](#)
- [Including Comments in ACLs, page 25-17](#)

Access List Numbers

The number you use to denote your ACL shows the type of access list that you are creating. [Table 25-1](#) lists the access-list number and corresponding access list type and shows whether or not they are supported in the switch. The switch supports IP standard and IP extended access lists, numbers 1 to 199 and 1300 to 2699.

Table 25-1 Access List Numbers

Access List 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 standard and extended named IP ACLs 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

Beginning in privileged EXEC mode, follow these steps to create a numbered standard ACL:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	access-list <i>access-list-number</i> { deny permit } <i>source</i> [<i>source-wildcard</i>] [log]	Define a standard IP access list by using a source address and wildcard. The <i>access-list-number</i> is a decimal number from 1 to 99 or 1300 to 1999. Enter deny or permit to specify whether to deny or permit access if conditions are matched. The <i>source</i> is the source address of the network or host from which the packet is being sent specified as: - 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. (Optional) The <i>source-wildcard</i> applies wildcard bits to the source. (Optional) Enter log to cause an informational logging message about the packet that matches the entry to be sent to the console.
Step 3	end	Return to privileged EXEC mode.
Step 4	show access-lists [<i>number</i> <i>name</i>]	Show the access list configuration.
Step 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
```

The switch always rewrites the order of standard access lists so that entries with **host** matches and entries with matches having a *don't care* mask of 0.0.0.0 are moved to the top of the list, above any entries with non-zero *don't care* masks. Therefore, in **show** command output and in the configuration file, the ACEs do not necessarily appear in the order in which they were entered.

The switch software can provide logging messages about packets permitted or denied by a standard IP access list. That is, any packet that matches the ACL causes an informational logging message about the packet to be sent to the console. The level of messages logged to the console is controlled by the logging console commands controlling the syslog messages.

**Note**

Because routing is done in hardware and logging is done in software, if a large number of packets match a *permit* or *deny* ACE containing a **log** keyword, the software might not be able to match the hardware processing rate, and not all packets will be logged.

The first packet that triggers the ACL causes a logging message right away, and subsequent packets are collected over 5-minute intervals before they are displayed or logged. The logging message includes the access list number, whether the packet was permitted or denied, the source IP address of the packet, and the number of packets from that source permitted or denied in the prior 5-minute interval.

After creating a numbered standard IP ACL, you can apply it to terminal lines (see the [“Applying an IP ACL to a Terminal Line” section on page 25-17](#)), to interfaces (see the [“Applying an IP ACL to an Interface” section on page 25-18](#)), or to VLANs (see the [“Configuring VLAN Maps” section on page 25-27](#)).

Creating a Numbered Extended ACL

Although standard ACLs use only source addresses for matching, you can use extended ACL source and destination addresses for matching operations and optional protocol type information for finer granularity of control. When you are creating ACEs in numbered extended access lists, remember that after you create the ACL, 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.

Some protocols also have specific parameters and keywords that apply to that protocol.

These IP protocols are supported (protocol keywords are in parentheses in bold):

Authentication Header Protocol (**ahp**), Enhanced Interior Gateway Routing Protocol (**eigrp**), Encapsulation Security Payload (**esp**), generic routing encapsulation (**gre**), Internet Control Message Protocol (**icmp**), Internet Group Management Protocol (**igmp**), Interior Gateway Routing Protocol (**igrp**), any Interior Protocol (**ip**), IP in IP tunneling (**ipinip**), KA9Q NOS-compatible IP over IP tunneling (**nos**), Open Shortest Path First routing (**ospf**), Payload Compression Protocol (**pcp**), Protocol Independent Multicast (**pim**), Transmission Control Protocol (**tcp**), or User Datagram Protocol (**udp**).

**Note**

ICMP echo-reply cannot be filtered. All other ICMP codes or types can be filtered.

For more details on the specific keywords relative to each protocol, refer to *Cisco IP and IP Routing Command Reference for IOS Release 12.1*.

**Note**

The switch does not support dynamic or reflexive access lists. It also does not support filtering based on the type of service (ToS) minimize-monetary-cost bit.

Supported parameters can be grouped into these categories: TCP, UDP, ICMP, IGMP, or other IP.

Beginning in privileged EXEC mode, follow these steps to create an extended ACL:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2a	access-list <i>access-list-number</i> { deny permit } <i>protocol</i> <i>source source-wildcard</i> <i>destination destination-wildcard</i> [precedence <i>precedence</i>] [tos <i>tos</i>] [fragments] [log] [log-input] [time-range <i>time-range-name</i>] [dscp <i>dscp</i>] Note If you enter a dscp value, you cannot enter tos or precedence . You can enter both a tos and a precedence value with no dscp .	Define an extended IP access list and the access conditions. The <i>access-list-number</i> is a decimal number from 100 to 199 or 2000 to 2699. Enter deny or permit to specify whether to deny or permit the packet if conditions are matched. For <i>protocol</i>, enter the name or number of an IP protocol: ahp, eigrp, esp, gre, icmp, igmp, igrp, ip, ipinip, nos, ospf, pcp, pim, tcp, or udp, or an integer in the range 0 to 255 representing an IP protocol number. To match any Internet protocol (including ICMP, TCP, and UDP) use the keyword ip. Note This step includes options for most IP protocols. For additional specific parameters for TCP, UDP, ICMP, and IGMP, see steps 2b through 2e. The <i>source</i> is the number of the network or host from which the packet is sent. The <i>source-wildcard</i> applies wildcard bits to the source. The <i>destination</i> is the network or host number to which the packet is sent. The <i>destination-wildcard</i> applies wildcard bits to the destination. Source, source-wildcard, destination, and destination-wildcard can be specified as: • The 32-bit quantity in dotted-decimal format. • The keyword any for 0.0.0.0 255.255.255.255 (any host). • The keyword host for a single host 0.0.0.0. The other keywords are optional and have these meanings: • precedence—Enter to match packets with a precedence level specified as a number from 0 to 7 or by name: routine (0), priority (1), immediate (2), flash (3), flash-override (4), critical (5), internet (6), network (7). • fragments—Enter to check non-initial fragments. • tos—Enter to match by type of service level, specified by a number from 0 to 15 or a name: normal (0), max-reliability (2), max-throughput (4), min-delay (8). • log—Enter to create an informational logging message to be sent to the console about the packet that matches the entry or log-input to include the input interface in the log entry. • time-range—For an explanation of this keyword, see the “Using Time Ranges with ACLs” section on page 25-15. • dscp—Enter to match packets with the DSCP value specified by a number from 0 to 63, or use the question mark (?) to see a list of available values.

	Command	Purpose
or	access-list <i>access-list-number</i> { deny permit } <i>protocol</i> any any [precedence <i>precedence</i>] [tos <i>tos</i>] [fragments] [log] [log-input] [time-range <i>time-range-name</i>] [dscp <i>dscp</i>]	In access-list configuration mode, define an extended IP access list using an abbreviation for a source and source wildcard of 0.0.0.0 255.255.255.255 and an abbreviation for a destination and destination wildcard of 0.0.0.0 255.255.255.255. You can use the any keyword in place of source and destination address and wildcard.
or	access-list <i>access-list-number</i> { deny permit } <i>protocol</i> host <i>source</i> host <i>destination</i> [precedence <i>precedence</i>] [tos <i>tos</i>] [fragments] [log] [log-input] [time-range <i>time-range-name</i>] [dscp <i>dscp</i>]	Define an extended IP access list using an abbreviation for a source and source wildcard of <i>source</i> 0.0.0.0 and an abbreviation for a destination and destination wildcard of <i>destination</i> 0.0.0.0. You can use the host keyword in place of source and destination wildcard or mask.
Step 2b	access-list <i>access-list-number</i> { deny permit } tcp <i>source</i> <i>source-wildcard</i> [<i>operator</i> <i>port</i>] <i>destination</i> <i>destination-wildcard</i> [<i>operator</i> <i>port</i>] [established] [precedence <i>precedence</i>] [tos <i>tos</i>] [fragments] [log] [log-input] [time-range <i>time-range-name</i>] [dscp <i>dscp</i>] [<i>flag</i>]	(Optional) Define an extended TCP access list and the access conditions. Enter tcp for Transmission Control Protocol. The parameters are the same as those described in Step 2a with these exceptions: (Optional) Enter an <i>operator</i> and <i>port</i> to compare source (if positioned after <i>source</i> <i>source-wildcard</i>) or destination (if positioned after <i>destination</i> <i>destination-wildcard</i>) port. Possible operators include eq (equal), gt (greater than), lt (less than), neq (not equal), and range (inclusive range). Operators require a port number (range requires two port numbers separated by a space). Enter the <i>port</i> number as a decimal number (from 0 to 65535) or the name of a TCP port. To see TCP port names, use the ? or refer to “Configuring IP Services” section of <i>Cisco IOS IP and IP Routing Command Reference for IOS Release 12.1</i> . Use only TCP port numbers or names when filtering TCP. The additional optional keywords have these meanings: • established—Enter to match an established connection. This has the same function as matching on the ack or rst flag. • <i>flag</i>—Enter one of these flags to match by the specified TCP header bits: ack (acknowledge), fin (finish), psh (push), rst (reset), syn (synchronize), or urg (urgent).
Step 2c	access-list <i>access-list-number</i> { deny permit } udp <i>source</i> <i>source-wildcard</i> [<i>operator</i> <i>port</i>] <i>destination</i> <i>destination-wildcard</i> [<i>operator</i> <i>port</i>] [precedence <i>precedence</i>] [tos <i>tos</i>] [fragments] [log] [log-input] [time-range <i>time-range-name</i>] [dscp <i>dscp</i>]	(Optional) Define an extended UDP access list and the access conditions. Enter udp for the User Datagram Protocol. The UDP parameters are the same as those described for TCP except that [<i>operator</i> [<i>port</i>]] port number or name must be a UDP port number or name, and the flag and established parameters are not valid for UDP.

	Command	Purpose
Step 2d	access-list <i>access-list-number</i> { deny permit } icmp <i>source source-wildcard destination destination-wildcard</i> [<i>icmp-type</i> [<i>icmp-type icmp-code</i>] [<i>icmp-message</i>]] [precedence <i>precedence</i>] [tos <i>tos</i>] [fragments] [log] [log-input] [time-range <i>time-range-name</i>] [dscp <i>dscp</i>]	(Optional) Define an extended ICMP access list and the access conditions. Enter icmp for Internet Control Message Protocol. The ICMP parameters are the same as those described for most IP protocols in Step 2a, with the addition of the ICMP message type and code parameters. These optional keywords have these meanings: <i>icmp-type</i>—Enter to filter by ICMP message type, a number from 0 to 255. <i>icmp-code</i>—Enter to filter ICMP packets that are filtered by ICMP message type by the ICMP message code, a number from 0 to 255. <i>icmp-message</i>—Enter to filter ICMP packets by ICMP message type name or ICMP message type and code name. To see a list of ICMP message type names and ICMP message type and code names, use the ? or refer to the “Configuring IP Services” section of <i>Cisco IOS IP and IP Routing Command Reference for IOS Release 12.1</i>.
Step 2e	access-list <i>access-list-number</i> { deny permit } igmp <i>source source-wildcard destination destination-wildcard</i> [<i>igmp-type</i>] [precedence <i>precedence</i>] [tos <i>tos</i>] [fragments] [log] [log-input] [time-range <i>time-range-name</i>] [dscp <i>dscp</i>]	(Optional) Define an extended IGMP access list and the access conditions. Enter igmp for Internet Group Management Protocol. The IGMP parameters are the same as those described for most IP protocols in Step 2a, with the addition of this optional parameter. <i>igmp-type</i>—To match IGMP message type, enter a number from 0 to 15, or enter the message name (dvmrp, host-query, host-report, pim, or trace).
Step 3	show access-lists [<i>number</i> <i>name</i>]	Verify the access list configuration.
Step 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 cannot selectively add or remove access list entries from a numbered access list.



Note

When you are creating an ACL, remember that, by default, the end of the access list contains an implicit deny statement for all packets if it did not find a match before reaching the end.

After creating a numbered extended ACL, you can apply it to terminal lines (see the [“Applying an IP ACL to a Terminal Line”](#) section on page 25-17), to interfaces (see the [“Applying an IP ACL to an Interface”](#) section on page 25-18), or to VLANs (see the [“Configuring VLAN Maps”](#) section on page 25-27).

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 in a router than if you were to 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 access list.



Note

The name you give to a standard 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:

- Not all commands that accept a numbered ACL accept a named ACL. ACLs for packet filters and route filters on interfaces can use a name. VLAN maps also accept a name.
- 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 25-6.
- You can use standard and extended ACLs (named or numbered) in VLAN maps.

Beginning in privileged EXEC mode, follow these steps to create a standard ACL using names:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip access-list standard <i>name</i>	Define a standard IP access list using a name, and enter access-list configuration mode. Note The name can be a number from 1 to 99.
Step 3	deny { <i>source</i> [<i>source-wildcard</i>] host <i>source</i> any } [log] or permit { <i>source</i> [<i>source-wildcard</i>] host <i>source</i> any } [log]	In access-list configuration mode, specify one or more conditions denied or permitted to determine if the packet is forwarded or dropped. • host <i>source</i>—A source and source wildcard of <i>source</i> 0.0.0.0. • any—A source and source wildcard of 0.0.0.0 255.255.255.255.
Step 4	end	Return to privileged EXEC mode.
Step 5	show access-lists [<i>number</i> <i>name</i>]	Show the access list configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a named standard ACL, use the **no ip access-list standard** *name* global configuration command.

Beginning in privileged EXEC mode, follow these steps to create an extended ACL using names:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip access-list extended <i>name</i>	Define an extended IP access list using a name and enter access-list configuration mode. Note The name can be a number from 100 to 199.
Step 3	{deny permit} protocol {source [source-wildcard] host source any} {destination [destination-wildcard] host destination any} [precedence precedence] [tos tos] [established] [log] [time-range time-range-name]	In access-list configuration mode, specify the conditions allowed or denied. Use the log keyword to get access list logging messages, including violations. See the “Creating a Numbered Extended ACL” section on page 25-9 for definitions of protocols and other keywords. • host source—A source and source wildcard of <i>source</i> 0.0.0.0. • host destination—A destination and destination wildcard of <i>destination</i> 0.0.0.0. • any—A source and source wildcard or destination and destination wildcard of 0.0.0.0 255.255.255.255.
Step 4	end	Return to privileged EXEC mode.
Step 5	show access-lists [<i>number</i> <i>name</i>]	Show the access list configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a named extended ACL, use the **no ip access-list extended** *name* global configuration command.

When you are creating standard extended ACLs, 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 ACL entries to a specific ACL. However, you can use **no permit** and **no deny** access-list configuration mode commands to remove entries from a named ACL. This example shows how you can delete individual ACEs from the named access list *border-list*:

```
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 a named ACL, you can apply it to interfaces (see the [“Applying an IP ACL to an Interface”](#) section on page 25-18) or VLANs (see the [“Configuring VLAN Maps”](#) section on page 25-27).

Using Time Ranges with ACLs

You can selectively apply extended ACLs based on the time of day and week by using the **time-range** global configuration command. First, define a time-range name and set the times and the dates or the days of the week in the time range. Then enter the time-range name when applying an ACL to set 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, for example, during a specified time period or on specified days of the week. The **time-range** keyword and argument are referenced in the named and numbered extended ACL task tables in the previous sections, the [“Creating Standard and Extended IP ACLs” section on page 25-6](#), and the [“Creating Named Standard and Extended ACLs” section on page 25-13](#).

These are some of the many possible 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 be set to log traffic only at certain times of the day. Therefore, you can simply deny access without needing to analyze many logs generated during peak hours.

Time-based access lists trigger CPU activity because the new configuration of the access list must be merged with other features and the combined configuration loaded into the TCAM. For this reason, you should be careful not to have several access lists configured to take affect in close succession (within a small number of minutes of each other.)



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 5-1](#).

Beginning in privileged EXEC mode, follow these steps to configure an time-range parameter for an ACL:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	time-range <i>time-range-name</i>	Assign a meaningful name (for example, <i>workhours</i>) to the time range to be created, and enter time-range configuration mode. The name cannot contain a space or quotation mark and must begin with a letter.
Step 3	absolute [<i>start time date</i>] [<i>end time date</i>] or periodic <i>day-of-the-week hh:mm to</i> <i>[day-of-the-week] 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. • You can use only one absolute statement in the time range. If you configure more than one absolute statement, only the one configured last is executed. • You can enter multiple periodic statements. For example, you could configure different hours for weekdays and weekends. Refer to the example configurations.
Step 4	end	Return to privileged EXEC mode.
Step 5	show time-range	Verify the time-range configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Repeat the steps if you have multiple items that you want in effect at different times.

To remove a configured time-range limitation, use the **no time-range** *time-range-name* global configuration command.

This example shows how to configure time ranges for *workhours* and for company holidays and 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_2003
Switch(config-time-range)# absolute start 00:00 1 Jan 2003 end 23:59 1 Jan 2003
Switch(config-time-range)# exit
Switch(config)# time-range thanksgiving_2003
Switch(config-time-range)# absolute start 00:00 27 Nov 2003 end 23:59 28 Nov 2003
Switch(config-time-range)# exit
Switch(config)# time-range christmas_2003
Switch(config-time-range)# absolute start 00:00 24 Dec 2003 end 23:59 25 Dec 2003
Switch(config-time-range)# end
Switch# show time-range
time-range entry: christmas_2003 (inactive)
    absolute start 00:00 24 December 2003 end 23:59 25 December 2003
time-range entry: new_year_day_2003 (inactive)
    absolute start 00:00 01 January 2003 end 23:59 01 January 2003
time-range entry: thanksgiving_2003 (inactive)
    absolute start 00:00 27 November 2003 end 23:59 28 November 2003
time-range entry: workhours (inactive)
    periodic weekdays 8:00 to 12:00
    periodic weekdays 13:00 to 17:00
```

To apply a time-range, enter the time-range name 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 times 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 in ACLs

You can use the **remark** keyword 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.

To include a comment for IP numbered standard or extended ACLs, use the **access-list access-list number remark remark** global configuration command. 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 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
```

Applying an IP ACL to a Terminal Line

You can use numbered ACLs to control access to one or more terminal lines. You cannot apply named ACLs to lines. You must set identical restrictions on all the virtual terminal lines because a user can attempt to connect to any of them.

For procedures for applying ACLs to interfaces, see the [“Applying an IP ACL to an Interface” section on page 25-18](#). For applying ACLs to VLANs, see the [“Configuring VLAN Maps” section on page 25-27](#).

Beginning in privileged EXEC mode, follow these steps to restrict incoming and outgoing connections between a virtual terminal line and the addresses in an ACL:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	line [console vty] line-number	Identify a specific line to configure, and enter in-line configuration mode. console—Specify the console terminal line. The console port is DCE. vty—Specify a virtual terminal for remote console access. The <i>line-number</i> is the first line number in a contiguous group that you want to configure when the line type is specified. The range is from 0 to 16.
Step 3	access-class access-list-number {in out}	Restrict incoming and outgoing connections between a particular virtual terminal line (into a device) and the addresses in an access list.

	Command	Purpose
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Display the access list configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove an ACL from a terminal line, use the **no access-class access-list-number {in | out}** line configuration command.

Applying an IP ACL to an Interface

This section describes how to apply IP ACLs to network interfaces. You can apply an ACL to *either* outbound or inbound Layer 3 interfaces. You can apply ACLs only to inbound Layer 2 interfaces. Note these guidelines:

- When controlling access to an interface, you can use a named or numbered ACL.
- If you apply an ACL to a Layer 2 interface that is a member of a VLAN, the Layer 2 (port) ACL takes precedence over an input Layer 3 ACL applied to the VLAN interface or a VLAN map applied to the VLAN. Incoming packets received on the Layer 2 port are always filtered by the port ACL.
- If you apply an ACL to a Layer 3 interface and routing is not enabled on the switch, the ACL only filters packets that are intended for the CPU, such as SNMP, Telnet, or web traffic. You do not have to enable routing to apply ACLs to Layer 2 interfaces.



Note

By default, the router sends Internet Control Message Protocol (ICMP) unreachable messages when a packet is denied by an access group. These access-group denied packets are not dropped in hardware but are bridged to the switch CPU so that it can generate the ICMP-unreachable message.

Beginning in privileged EXEC mode, follow these steps to control access to an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface interface-id	Identify a specific interface for configuration, and enter interface configuration mode. The interface can be a Layer 2 interface (port ACL), or a Layer 3 interface (router ACL).
Step 3	ip access-group {access-list-number name} {in out}	Control access to the specified interface. The out keyword is not supported for Layer 2 interfaces (port ACLs).
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Display the access list configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the specified access group, use the **no ip access-group {access-list-number | name} {in | out}** interface configuration command.

This example shows how to apply access list 2 on an interface to filter packets entering the interface:

```
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# ip access-group 2 in
```

**Note**

When you apply the **ip access-group** interface configuration command to a Layer 3 interface (an SVI, a Layer 3 EtherChannel, or a routed port), the interface must have been configured with an IP address. Layer 3 access groups filter packets that are routed or are received by Layer 3 processes on the CPU. They do not affect packets bridged within a VLAN.

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.

For outbound ACLs, after receiving and routing a packet to a controlled interface, the switch checks the packet against the ACL. If the ACL permits the packet, the switch sends the packet. If the ACL rejects the packet, the switch discards the packet.

By default, the input interface sends ICMP Unreachable messages whenever a packet is discarded, regardless of whether the packet was discarded because of an ACL on the input interface or because of an ACL on the output interface. ICMP Unreachables are normally limited to no more than one every one-half second per input interface, but this can be changed by using the **ip icmp rate-limit unreachable** global configuration command.

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.

Hardware and Software Treatment of IP ACLs

ACL processing is primarily accomplished in hardware, but requires forwarding of some traffic flows to the CPU for software processing. If the hardware reaches its capacity to store ACL configurations, packets are sent to the CPU for forwarding. The forwarding rate for software-forwarded traffic is substantially less than for hardware-forwarded traffic.

For router ACLs, other factors can cause packets to be sent to the CPU:

- Using the **log** keyword
- Generating ICMP unreachable messages

When traffic flows are both logged and forwarded, forwarding is done by hardware, but logging must be done by software. Because of the difference in packet handling capacity between hardware and software, if the sum of all flows being logged (both permitted flows and denied flows) is of great enough bandwidth, not all of the packets that are forwarded can be logged.

If router ACL configuration cannot be applied in hardware, packets arriving in a VLAN that must be routed are routed in software, but are bridged in hardware. If ACLs cause large numbers of packets to be sent to the CPU, the switch performance can be negatively affected.

When you enter the **show ip access-lists** privileged EXEC command, the match count displayed does not account for packets that are access controlled in hardware. Use the **show access-lists hardware counters** privileged EXEC command to obtain some basic hardware ACL statistics for switched and routed packets.

Router ACLs function as follows:

- The hardware controls permit and deny actions of standard and extended ACLs (input and output) for security access control.
- If **log** has not been specified, the flows that match a *deny* statement in a security ACL are dropped by the hardware if *ip unreachable* is disabled. The flows matching a *permit* statement are switched in hardware.
- Adding the **log** keyword to an ACE in a router ACL causes a copy of the packet to be sent to the CPU for logging only. If the ACE is a *permit* statement, the packet is still switched and routed in hardware.

IP ACL Configuration Examples


This section provides examples of configuring and applying IP ACLs. For detailed information about compiling ACLs, refer to the *Security Configuration Guide* and the “IP Services” chapter of the *Cisco IOS IP and IP Routing Configuration Guide for IOS Release 12.1*.

Figure 25-3 shows a small networked office environment with the routed Port 2 connected to Server A, containing benefits and other information that all employees can access, and routed Port 1 connected to Server B, containing confidential payroll data. All users can access Server A, but Server B has restricted access.

Use router ACLs to do this in one of two ways:

- Create a standard ACL, and filter traffic coming to the server from Port 1.
- Create an extended ACL, and filter traffic coming from the server into Port 1.

Figure 25-3 Using Router ACLs to Control Traffic



This example uses a standard ACL to filter traffic coming into Server B from an interface, permitting traffic only from Accounting's source addresses 172.20.128.64 to 172.20.128.95. The ACL is applied to traffic coming out of routed Port 1 from the specified source address.

```
Switch(config)# access-list 6 permit 172.20.128.64 0.0.0.31
Switch(config)# end
Switch# show access-lists
Standard IP access list 6
    permit 172.20.128.64, wildcard bits 0.0.0.31
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# ip access-group 6 out
```

This example uses an extended ACL to filter traffic coming from Server B into Port 1, permitting traffic from any source address (in this case Server B) to only the Accounting destination addresses 172.20.128.64 to 172.20.128.95. The ACL is applied to traffic going into routed Port 1, permitting it to go only to the specified destination addresses. Note that with extended ACLs, you must enter the protocol (IP) before the source and destination information.

```
Switch(config)# access-list 106 permit ip any 172.20.128.64 0.0.0.31
Switch(config)# end
Switch# show access-lists
Extended IP access list 106
    permit ip any 172.20.128.64 0.0.0.31
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# ip access-group 106 in
```

Numbered ACLs

In this example, network 36.0.0.0 is a Class A network whose second octet specifies a subnet; that is, its subnet mask is 255.255.0.0. The third and fourth octets of a network 36.0.0.0 address specify a particular host. Using access list 2, the switch accepts one address on subnet 48 and reject all others on that subnet. The last line of the list shows that the switch accepts addresses on all other network 36.0.0.0 subnets. The ACL is then applied to packets entering an interface.

```
Switch(config)# access-list 2 permit 36.48.0.3
Switch(config)# access-list 2 deny 36.48.0.0 0.0.255.255
Switch(config)# access-list 2 permit 36.0.0.0 0.255.255.255
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# ip access-group 2 in
```

Extended ACLs

In this example, the first line permits any incoming TCP connections with destination ports greater than 1023. The second line permits incoming TCP connections to the Simple Mail Transfer Protocol (SMTP) port of host 128.88.1.2. The third line permits incoming ICMP messages for error feedback.

```
Switch(config)# access-list 102 permit tcp any 128.88.0.0 0.0.255.255 gt 1023
Switch(config)# access-list 102 permit tcp any host 128.88.1.2 eq 25
Switch(config)# access-list 102 permit icmp any any
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# ip access-group 102 in
```

For another example of using an extended ACL, suppose that you have a network connected to the Internet, and you want any host on the network to be able to form TCP connections to any host on the Internet. However, you do not want IP hosts to be able to form TCP connections to hosts on your network, except to the mail (SMTP) port of a dedicated mail host.

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. Outbound packets have the port numbers reversed. Because the secure system of the network always accepts mail connections on port 25, the incoming and outgoing services are separately controlled. The ACL must be configured as an input ACL on the outbound interface and an output ACL on the inbound interface.

In this example, the network is a Class B network with the address 128.88.0.0, and the mail host address is 128.88.1.2. The **established** keyword is used only for the TCP to show an established connection. A match occurs if the TCP datagram has the ACK or RST bits set, which show that the packet belongs to an existing connection. The specified interface is the interface that connects the router to the Internet.

```
Switch(config)# access-list 102 permit tcp any 128.88.0.0 0.0.255.255 established
Switch(config)# access-list 102 permit tcp any host 128.88.1.2 eq 25
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# ip access-group 102 in
```

Named ACLs

This example creates a standard ACL named *internet_filter* and an extended ACL named *marketing_group*. The *internet_filter* ACL allows all traffic from the source address 1.2.3.4.

```
Switch(config)# ip access-list standard Internet_filter
Switch(config-ext-nacl)# permit 1.2.3.4
Switch(config-ext-nacl)# exit
```

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 ICMP traffic, denies UDP traffic from any source to the destination address range 171.69.0.0 through 179.69.255.255 with a destination port less than 1024, denies any other IP traffic, and provides a log of the result.

```
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 icmp any any
Switch(config-ext-nacl)# deny udp any 171.69.0.0 0.0.255.255 lt 1024
Switch(config-ext-nacl)# deny ip any any log
Switch(config-ext-nacl)# exit
```

The ACLs are applied to an interface, which is configured as a Layer 3 port, with the *Internet_filter* ACL applied to outgoing traffic and the *marketing_group* ACL applied to incoming traffic.

```
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# no switchport
Switch(config-if)# ip address 2.0.5.1 255.255.255.0
Switch(config-if)# ip access-group Internet_filter out
Switch(config-if)# ip access-group marketing_group in
```

Time Range Applied to an IP ACL

This example denies HTTP traffic on IP on Monday through Friday between the hours of 8:00 a.m. and 6:00 p.m (18:00). The example allows UDP traffic only on Saturday and Sunday from noon to 8:00 p.m. (20:00).

```
Switch(config)# time-range no-http
Switch(config)# periodic weekdays 8:00 to 18:00
!
Switch(config)# time-range udp-yes
Switch(config)# periodic weekend 12:00 to 20:00
!
Switch(config)# ip access-list extended strict
Switch(config-ext-nacl)# deny tcp any any eq www time-range no-http
Switch(config-ext-nacl)# permit udp any any time-range udp-yes
!
Switch(config-ext-nacl)# exit
Switch(config)# interface fastethernet1/0/1
Switch(config-if)# ip access-group strict in
```

Commented IP ACL Entries

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
```

ACL Logging

Two variations of logging are supported on router ACLs. The **log** keyword sends an informational logging message to the console about the packet that matches the entry; the **log-input** keyword includes the input interface in the log entry.

In this example, standard named access list *stan1* denies traffic from 10.1.1.0 0.0.0.255, allows traffic from all other sources, and includes the **log** keyword.

```
Switch(config)# ip access-list standard stan1
Switch(config-std-nacl)# deny 10.1.1.0 0.0.0.255 log
Switch(config-std-nacl)# permit any log
Switch(config-std-nacl)# exit
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# ip access-group stan1 in
Switch(config-if)# end
Switch# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Console logging: level debugging, 37 messages logged
  Monitor logging: level debugging, 0 messages logged
  Buffer logging: level debugging, 37 messages logged
  File logging: disabled
  Trap logging: level debugging, 39 message lines logged
```

Log Buffer (4096 bytes):

00:00:48: NTP: authentication delay calculation problems

<output truncated>

```
00:09:34:%SEC-6-IPACCESSLOGS:list stan1 permitted 0.0.0.0 1 packet
00:09:59:%SEC-6-IPACCESSLOGS:list stan1 denied 10.1.1.15 1 packet
00:10:11:%SEC-6-IPACCESSLOGS:list stan1 permitted 0.0.0.0 1 packet
```

This example is a named extended access list *ext1* that permits ICMP packets from any source to 10.1.1.0 0.0.0.255 and denies all UDP packets.

```
Switch(config)# ip access-list extended ext1
Switch(config-ext-nacl)# permit icmp any 10.1.1.0 0.0.0.255 log
Switch(config-ext-nacl)# deny udp any any log
Switch(config-ext-nacl)# exit
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# ip access-group ext1 in
```

This is an example of a log for an extended ACL:

```
01:24:23:%SEC-6-IPACCESSLOGDP:list ext1 permitted icmp 10.1.1.15 -> 10.1.1.61 (0/0), 1
packet
01:25:14:%SEC-6-IPACCESSLOGDP:list ext1 permitted icmp 10.1.1.15 -> 10.1.1.61 (0/0), 7
packets
01:26:12:%SEC-6-IPACCESSLOGDP:list ext1 denied udp 0.0.0.0(0) -> 255.255.255.255(0), 1
packet
01:31:33:%SEC-6-IPACCESSLOGDP:list ext1 denied udp 0.0.0.0(0) -> 255.255.255.255(0), 8
packets
```

Note that all logging entries for IP ACLs start with %SEC-6-IPACCESSLOG with minor variations in format depending on the kind of ACL and the access entry that has been matched.

This is an example of an output message when the **log-input** keyword is entered:

```
00:04:21:%SEC-6-IPACCESSLOGDP:list inputlog permitted icmp 10.1.1.10 (Vlan1
0001.42ef.a400) -> 10.1.1.61 (0/0), 1 packet
```

A log message for the same sort of packet using the **log** keyword does not include the input interface information:

```
00:05:47:%SEC-6-IPACCESSLOGDP:list inputlog permitted icmp 10.1.1.10 -> 10.1.1.61 (0/0), 1 packet
```

Creating Named MAC Extended ACLs

You can filter non-IPv4 traffic on a VLAN or on a Layer 2 interface by using MAC addresses and named MAC extended ACLs. The procedure is similar to that of configuring other extended named ACLs.



Note

You cannot apply named MAC extended ACLs to Layer 3 interfaces.

For more information about the supported non-IPv4 protocols in the **mac access-list extended** command, refer to the command reference for this release.



Note

Though visible in the command-line help strings, **appletalk** is not supported as a matching condition for the **deny** and **permit** MAC access-list configuration mode commands.

Beginning in privileged EXEC mode, follow these steps to create a named MAC extended ACL:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mac access-list extended <i>name</i>	Define an extended MAC access list using a name.
Step 3	{deny permit} {any host <i>source MAC address</i> <i>source MAC address mask</i> } {any host <i>destination MAC address</i> <i>destination MAC address mask</i> [type <i>mask</i> lsap <i>lsap mask</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 0-65535] [cos <i>cos</i>]	In extended MAC access-list configuration mode, specify to permit or deny any source MAC address, a source MAC address with a mask, or a specific host source MAC address and any destination MAC address, destination MAC address with a mask, or a specific destination MAC address. (Optional) You can also enter these options: type mask—An arbitrary EtherType number of a packet with Ethernet II or SNAP encapsulation in decimal, hex, or octal with optional mask of <i>don't care</i> bits applied to the EtherType before testing for a match. lsap lsap mask—An LSAP number of a packet with 802.2 encapsulation in decimal, hex, or octal with optional mask of <i>don't care</i> bits. 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—A non-IP protocol. cos cos—An IEEE 802.1p cost of service number from 0 to 7 used to set priority.
Step 4	end	Return to privileged EXEC mode.

	Command	Purpose
Step 5	show access-lists [<i>number</i> <i>name</i>]	Show the access list configuration.
Step 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-lists
Extended MAC access list mac1
    deny any any decnet-iv
    permit any any
```

Applying a MAC ACL to a Layer 2 Interface

After you create a MAC ACL, you can apply it to a Layer 2 interface to filter non-IPv4 traffic coming in that interface. When you apply the MAC ACL, consider these guidelines:

- If you apply an ACL to a Layer 2 interface that is a member of a VLAN, the Layer 2 (port) ACL takes precedence over an input Layer 3 ACL applied to the VLAN interface or a VLAN map applied to the VLAN. Incoming packets received on the Layer 2 port are always filtered by the port ACL.
- You can apply no more than one IP access list and one MAC access list to the same Layer 2 interface. The IP access list filters only IPv4 packets, and the MAC access list filters non-IPv4 packets.
- A Layer 2 interface can have only one MAC access list. If you apply a MAC access list to a Layer 2 interface that has a MAC ACL configured, the new ACL replaces the previously configured one.

Beginning in privileged EXEC mode, follow these steps to apply a MAC access list to control access to a Layer 2 interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Identify a specific interface, and enter interface configuration mode. The interface must be a physical Layer 2 interface (port ACL).
Step 3	mac access-group { <i>name</i> } { in }	Control access to the specified interface by using the MAC access list. Note Port ACLs are supported only in the inbound direction.
Step 4	end	Return to privileged EXEC mode.
Step 5	show mac access-group [interface <i>interface-id</i>]	Display the MAC access list applied to the interface or all Layer 2 interfaces.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the specified access group, use the **no mac access-group** {name} interface configuration command.

This example shows how to apply MAC access list *mac1* on an interface to filter packets entering the interface:

```
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# mac access-group mac1 in
```

**Note**

The **mac access-group** interface configuration command is only valid when applied to a physical Layer 2 interface. You cannot use the command on EtherChannel port channels.

After receiving a packet, the switch checks it against the inbound ACL. If the ACL permits it, the switch continues to process the packet. If the ACL rejects the packet, the switch discards it. When you apply an undefined ACL to an interface, the switch acts as if the ACL has not been applied and permits all packets. Remember this behavior if you use undefined ACLs for network security.

Configuring VLAN Maps

This section describes how to configure VLAN maps, which is the only way to control filtering within a VLAN. VLAN maps have no direction. To filter traffic in a specific direction by using a VLAN map, you need to include an ACL with specific source or destination addresses. If there is a match clause for that type of packet (IP or MAC) in the VLAN map, the default action is to drop the packet if the packet does not match any of the entries within the map. If there is no match clause for that type of packet, the default is to forward the packet.

**Note**

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

To create a VLAN map and apply it to one or more VLANs, perform these steps:

- Step 1** Create the standard or extended IP ACLs or named MAC extended ACLs that you want to apply to the VLAN. See the [“Creating Standard and Extended IP ACLs”](#) section on page 25-6 and the [“Creating a VLAN Map”](#) section on page 25-28.
- Step 2** Enter the **vlan access-map** global configuration command to create a VLAN ACL map entry.
- Step 3** In access map configuration mode, optionally enter an **action—forward** (the default) or **drop**—and enter the **match** command to specify an IPv4 packet or a non-IPv4 packet (with only a known MAC address) and to match the packet against one or more ACLs (standard or extended).

**Note**

If the VLAN map has a match clause for the type of packet (IP or MAC) and the packet does not match the type, the default is to drop the packet. If there is no match clause in the VLAN map for that type of packet, and no action specified, the packet is forwarded.

- Step 4** Use the **vlan filter** global configuration command to apply a VLAN map to one or more VLANs.

This section contains these topics:

- [VLAN Map Configuration Guidelines, page 25-28](#)
- [Creating a VLAN Map, page 25-28](#)
- [Applying a VLAN Map to a VLAN, page 25-31](#)
- [Using VLAN Maps in Your Network, page 25-31](#)

VLAN Map Configuration Guidelines

Follow these guidelines when configuring VLAN maps:

- If there is no ACL configured to deny traffic on an interface and *no* VLAN map is configured, all traffic is permitted.
- Each VLAN map consists of a series of entries. The order of entries in an VLAN map is important. A packet that comes into the switch is tested against the first entry in the VLAN map. If it matches, the action specified for that part of the VLAN map is taken. If there is no match, the packet is tested against the next entry in the map.
- If the VLAN map has at least one match clause for the type of packet (IP or MAC) and the packet does not match any of these match clauses, the default is to drop the packet. If there is no match clause for that type of packet in the VLAN map, the default is to forward the packet.
- The system might take longer to boot if you have configured a very large number of ACLs.
- Logging is not supported for VLAN maps.
- If VLAN map configuration cannot be applied in hardware, all packets in that VLAN must be bridged and routed by software.
- When a switch has an IP access list or MAC access list applied to a Layer 2 interface, and you apply a VLAN map to a VLAN that the port belongs to, the port ACL takes precedence over the VLAN map.
- See the [“Using VLAN Maps in Your Network” section on page 25-31](#) for configuration examples.
- For information about using both router ACLs and VLAN maps, see the [“Guidelines” section on page 25-34](#).

Creating a VLAN Map

Each VLAN map consists of an ordered series of entries. Beginning in privileged EXEC mode, follow these steps to create, add to, or delete a VLAN map entry:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vlan access-map <i>name</i> [<i>number</i>]	Create a VLAN map, and give it a name and (optionally) a number. The number is the sequence number of the entry within the map. When you create VLAN maps with the same name, numbers are assigned sequentially in increments of 10. When modifying or deleting maps, you can enter the number of the map entry that you want to modify or delete. Entering this command changes to access-map configuration mode.

	Command	Purpose
Step 3	action {drop forward}	(Optional) Set the action for the map entry. The default is to forward.
Step 4	match {ip mac} address {name number} [name number]	Match the packet (using either the IP or MAC address) against one or more standard or extended access lists. Note that packets are only matched against access lists of the correct protocol type. IP packets are matched against standard or extended IP access lists. Non-IPv4 packets are only matched against named MAC extended access lists.
Step 5	end	Return to global configuration mode.
Step 6	show running-config	Display the access list configuration.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no vlan access-map name** global configuration command to delete a map.

Use the **no vlan access-map name number** global configuration command to delete a single sequence entry from within the map.

Use the **no action** access-map configuration command to enforce the default action, which is to forward.

VLAN maps do not use the specific permit or deny keywords. To deny a packet by using VLAN maps, create an ACL that would match the packet, and set the action to drop. A permit in the ACL counts as a match. A deny in the ACL means no match.

Examples of ACLs and VLAN Maps

These examples show how to create ACLs and VLAN maps that for specific purposes.

Example 1

This example shows how to create an ACL and a VLAN map to deny a packet. In the first map, any packets that match the *ip1* ACL (TCP packets) would be dropped. You first create the *ip1* ACL to permit any TCP packet and no other packets. Because there is a match clause for IP packets in the VLAN map, the default action is to drop any IP packet that does not match any of the match clauses.

```
Switch(config)# ip access-list extended ip1
Switch(config-ext-nacl)# permit tcp any any
Switch(config-ext-nacl)# exit
Switch(config)# vlan access-map map_1 10
Switch(config-access-map)# match ip address ip1
Switch(config-access-map)# action drop
```

This example shows how to create a VLAN map to permit a packet. ACL *ip2* permits UDP packets and any packets that match the *ip2* ACL are forwarded. In this map, any IP packets that did not match any of the previous ACLs (that is, packets that are not TCP packets or UDP packets) would get dropped.

```
Switch(config)# ip access-list extended ip2
Switch(config-ext-nacl)# permit udp any any
Switch(config-ext-nacl)# exit
Switch(config)# vlan access-map map_1 20
Switch(config-access-map)# match ip address ip2
Switch(config-access-map)# action forward
```

Example 2

In this example, the VLAN map has a default action of drop for IP packets and a default action of forward for MAC packets. Used with standard ACL 101 and extended named access lists **igmp-match** and **tcp-match**, the map will have the following results:

- Forward all UDP packets
- Drop all IGMP packets
- Forward all TCP packets
- Drop all other IP packets
- Forward all non-IPv4 packets

```
Switch(config)# access-list 101 permit udp any any
Switch(config)# ip access-list extended igmp-match
Switch(config-ext-nacl)# permit igmp any any
Switch(config)# ip access-list extended tcp-match
Switch(config-ext-nacl)# permit tcp any any
Switch(config-ext-nacl)# exit
Switch(config)# vlan access-map drop-ip-default 10
Switch(config-access-map)# match ip address 101
Switch(config-access-map)# action forward
Switch(config-access-map)# exit
Switch(config)# vlan access-map drop-ip-default 20
Switch(config-access-map)# match ip address igmp-match
Switch(config-access-map)# action drop
Switch(config-access-map)# exit
Switch(config)# vlan access-map drop-ip-default 30
Switch(config-access-map)# match ip address tcp-match
Switch(config-access-map)# action forward
```

Example 3

In this example, the VLAN map has a default action of drop for MAC packets and a default action of forward for IP packets. Used with MAC extended access lists **good-hosts** and **good-protocols**, the map will have the following results:

- Forward MAC packets from hosts 0000.0c00.0111 and 0000.0c00.0211
- Forward MAC packets with decnet-iv or vines-ip protocols
- Drop all other non-IPv4 packets
- Forward all IP packets

```
Switch(config)# mac access-list extended good-hosts
Switch(config-ext-macl)# permit host 000.0c00.0111 any
Switch(config-ext-macl)# permit host 000.0c00.0211 any
Switch(config-ext-nacl)# exit
Switch(config)# mac access-list extended good-protocols
Switch(config-ext-macl)# permit any any decnet-ip
Switch(config-ext-macl)# permit any any vines-ip
Switch(config-ext-nacl)# exit
Switch(config)# vlan access-map drop-mac-default 10
Switch(config-access-map)# match mac address good-hosts
Switch(config-access-map)# action forward
Switch(config-access-map)# exit
Switch(config)# vlan access-map drop-mac-default 20
Switch(config-access-map)# match mac address good-protocols
Switch(config-access-map)# action forward
```

Example 4

In this example, the VLAN map has a default action of drop for all packets (IPv4 and non-IPv4). Used with access lists **tcp-match** and **good-hosts** from Examples 2 and 3, the map will have the following results:

- Forward all TCP packets
- Forward MAC packets from hosts 0000.0c00.0111 and 0000.0c00.0211
- Drop all other IP packets
- Drop all other MAC packets

```
Switch(config)# vlan access-map drop-all-default 10
Switch(config-access-map)# match ip address tcp-match
Switch(config-access-map)# action forward
Switch(config-access-map)# exit
Switch(config)# vlan access-map drop-all-default 20
Switch(config-access-map)# match mac address good-hosts
Switch(config-access-map)# action forward
```

Applying a VLAN Map to a VLAN

Beginning in privileged EXEC mode, follow these steps to apply a VLAN map to one or more VLANs:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vlan filter <i>mapname</i> vlan-list <i>list</i>	Apply the VLAN map to one or more VLAN IDs. The list can be a single VLAN ID (22), a consecutive list (10-22), or a string of VLAN IDs (12, 22, 30). Spaces around the comma and hyphen are optional.
Step 3	show running-config	Display the access list configuration.
Step 4	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the VLAN map, use the **no vlan filter** *mapname* **vlan-list** *list* global configuration command.

This example shows how to apply VLAN map 1 to VLANs 20 through 22:

```
Switch(config)# vlan filter map 1 vlan-list 20-22
```

Using VLAN Maps in Your Network


This section describes some typical uses for VLAN maps and includes these topics:

- [Wiring Closet Configuration, page 25-32](#)
- [Denying Access to a Server on Another VLAN, page 25-33](#)

Wiring Closet Configuration

In a wiring closet configuration, routing might not be enabled on the switch. In this configuration, the switch can still support a VLAN map and a QoS classification ACL. In [Figure 25-4](#), assume that Host X and Host Y are in different VLANs and are connected to wiring closet switches A and C. Traffic from Host X to Host Y is eventually being routed by Switch B, which has routing enabled. Traffic from Host X to Host Y can be access-controlled at the traffic entry point, Switch A.

Figure 25-4 Wiring Closet Configuration



If you do not want HTTP traffic switched from Host X to Host Y, you can configure a VLAN map on Switch A to drop all HTTP traffic from Host X (IP address 10.1.1.32) to Host Y (IP address 10.1.1.34) at Switch A and not bridge it to Switch B.

First, define the IP access list *http* that permits (matches) any TCP traffic on the HTTP port.

```
Switch(config)# ip access-list extended http
Switch(config-ext-nacl)# permit tcp host 10.1.1.32 host 10.1.1.34 eq www
Switch(config-ext-nacl)# exit
```

Next, create VLAN access map *map2* so that traffic that matches the *http* access list is dropped and all other IP traffic is forwarded.

```
Switch(config)# vlan access-map map2 10
Switch(config-access-map)# match ip address http
Switch(config-access-map)# action drop
Switch(config-access-map)# exit
Switch(config)# ip access-list extended match_all
Switch(config-ext-nacl)# permit ip any any
Switch(config-ext-nacl)# exit
Switch(config)# vlan access-map map2 20
Switch(config-access-map)# match ip address match_all
Switch(config-access-map)# action forward
```

Then, apply VLAN access map *map2* to VLAN 1.


```
Switch(config)# vlan filter map2 vlan 1
```

Denying Access to a Server on Another VLAN

You can restrict access to a server on another VLAN. For example, server 10.1.1.100 in VLAN 10 needs to have access denied to these hosts (see Figure 25-5).

- Hosts in subnet 10.1.2.0/8 in VLAN 20 should not have access.
- Hosts 10.1.1.4 and 10.1.1.8 in VLAN 10 should not have access.

Figure 25-5 Deny Access to a Server on Another VLAN



This example shows how to deny access to a server on another VLAN by creating the VLAN map SERVER1 that denies access to hosts in subnet 10.1.2.0/8, host 10.1.1.4, and host 10.1.1.8 and permits other IP traffic. The final step is to apply the map SERVER1 to VLAN 10.

Step 1 Define the IP ACL that will match the correct packets.

```
Switch(config)# ip access-list extended SERVER1_ACL
Switch(config-ext-nacl)# permit ip 10.1.2.0 0.0.0.255 host 10.1.1.100
Switch(config-ext-nacl)# permit ip host 10.1.1.4 host 10.1.1.100
Switch(config-ext-nacl)# permit ip host 10.1.1.8 host 10.1.1.100
Switch(config-ext-nacl)# exit
```

Step 2 Define a VLAN map using this ACL that will drop IP packets that match SERVER1_ACL and forward IP packets that do not match the ACL.

```
Switch(config)# vlan access-map SERVER1_MAP
Switch(config-access-map)# match ip address SERVER1_ACL
Switch(config-access-map)# action drop
Switch(config)# vlan access-map SERVER1_MAP 20
Switch(config-access-map)# action forward
Switch(config-access-map)# exit
```

Step 3 Apply the VLAN map to VLAN 10.

```
Switch(config)# vlan filter SERVER1_MAP vlan-list 10.
```

Using VLAN Maps with Router ACLs

To access control both bridged and routed traffic, you can use VLAN maps only or a combination of router ACLs and VLAN maps. You can define router ACLs on both input and output routed VLAN interfaces, and you can define a VLAN map to access control the bridged traffic.

If a packet flow matches a VLAN-map deny clause in the ACL, regardless of the router ACL configuration, the packet flow is denied.

**Note**

When you use router ACLs with VLAN maps, packets that require logging on the router ACLs are not logged if they are denied by a VLAN map.

If the VLAN map has a match clause for the type of packet (IP or MAC) and the packet does not match the type, the default is to drop the packet. If there is no match clause in the VLAN map, and no action specified, the packet is forwarded if it does not match any VLAN map entry.

This section includes this information about using VLAN maps with router ACLs:

- [Guidelines, page 25-34](#)
- [Examples of Router ACLs and VLAN Maps Applied to VLANs, page 25-35](#)

Guidelines

These guidelines are for configurations where you need to have an router ACL *and* a VLAN map on the same VLAN. These guidelines do not apply to configurations where you are mapping router ACLs and VLAN maps on different VLANs.

The switch hardware provides one lookup for security ACLs for each direction (input and output); therefore, you must merge a router ACL and a VLAN map when they are configured on the same VLAN. Merging the router ACL with the VLAN map might significantly increase the number of ACEs.

If you must configure a router ACL and a VLAN map on the same VLAN, use these guidelines for both router ACL and VLAN map configuration:

- You can configure only one VLAN map and one router ACL in each direction (input/output) on a VLAN interface.
- Whenever possible, try to write the ACL with all entries having a single action except for the final, default action of the other type. That is, write the ACL using one of these two forms:

permit...

permit...

permit...

deny ip any any

or

deny...

deny...

deny...

permit ip any any

- To define multiple actions in an ACL (permit, deny), group each action type together to reduce the number of entries.

- Avoid including Layer 4 information in an ACL; adding this information complicates the merging process. The best merge results are obtained if the ACLs are filtered based on IP addresses (source and destination) and not on the full flow (source IP address, destination IP address, protocol, and protocol ports). It is also helpful to use *don't care* bits in the IP address, whenever possible.

If you need to specify the full-flow mode and the ACL contains both IP ACEs and TCP/UDP/ICMP ACEs with Layer 4 information, put the Layer 4 ACEs at the end of the list. This gives priority to the filtering of traffic based on IP addresses.


Examples of Router ACLs and VLAN Maps Applied to VLANs

This section gives examples of applying router ACLs and VLAN maps to a VLAN for switched, bridged, routed, and multicast packets. Although the following illustrations show packets being forwarded to their destination, each time the packet's path crosses a line indicating a VLAN map or an ACL, it is also possible that the packet might be dropped, rather than forwarded.

ACLs and Switched Packets

Figure 25-6 shows how an ACL is applied on packets that are switched within a VLAN. Packets switched within the VLAN without being routed or forwarded by fallback bridging are only subject to the VLAN map of the input VLAN.


Figure 25-6 Applying ACLs on Switched Packets



ACLs and Bridged Packets

Figure 25-7 shows how an ACL is applied on fallback-bridged packets. For bridged packets, only Layer 2 ACLs are applied to the input VLAN. Only non-IPv4, non-ARP packets can be fallback-bridged.

Figure 25-7 Applying ACLs on Bridged Packets




ACLs and Routed Packets

Figure 25-8 shows how ACLs are applied on routed packets. For routed packets, the ACLs are applied in this order:

1. VLAN map for input VLAN
2. Input router ACL
3. Output router ACL
4. VLAN map for output VLAN

Figure 25-8 Applying ACLs on Routed Packets




ACLs and Multicast Packets

Figure 25-9 shows how ACLs are applied on packets that are replicated for IP multicasting. A multicast packet being routed has two different kinds of filters applied: one for destinations that are other ports in the input VLAN and another for each of the destinations that are in other VLANs to which the packet has been routed. The packet might be routed to more than one output VLAN, in which case a different router output ACL and VLAN map would apply for each destination VLAN.

The final result is that the packet might be permitted in some of the output VLANs and not in others. A copy of the packet is forwarded to those destinations where it is permitted. However, if the input VLAN map (VLAN 10 map in Figure 25-9) drops the packet, no destination receives a copy of the packet.

Figure 25-9 Applying ACLs on Multicast Packets



Displaying ACL Configuration

You can display the ACLs that are configured on the switch, and you can display the ACLs that have been applied to interfaces and VLANs.

When you use the **ip access-group** interface configuration command to apply ACLs to a Layer 2 or 3 interface, you can display the access groups on the interface. You can also display the MAC ACLs applied to a Layer 2 interface. You can use the privileged EXEC commands as described in [Table 25-2](#) to display this information.

Table 25-2 Commands for Displaying Access Lists and Access Groups

Command	Purpose
show access-lists [<i>number</i> <i>name</i>]	Display the contents of one or all current IP and MAC address access lists or a specific access list (numbered or named).
show ip access-lists [<i>number</i> <i>name</i>]	Display the contents of all current IP access lists or a specific IP access list (numbered or named).
show ip interface <i>interface-id</i>	Display detailed configuration and status of an interface. If IP is enabled on the interface and ACLs have been applied by using the ip access-group interface configuration command, the access groups are included in the display.
show running-config [interface <i>interface-id</i>]	Displays the contents of the configuration file for the switch or the specified interface, including all configured MAC and IP access lists and which access groups are applied to an interface.
show mac access-group [interface <i>interface-id</i>]	Displays MAC access lists applied to all Layer 2 interfaces or the specified Layer 2 interface.

You can also display information about VLAN access maps or VLAN filters. Use the privileged EXEC commands in [Table 25-3](#) to display VLAN map information.

Table 25-3 Commands for Displaying VLAN Map Information

Command	Purpose
show vlan access-map [<i>mapname</i>]	Show information about all VLAN access-maps or the specified access map.
show vlan filter [access-map <i>name</i> vlan <i>vlan-id</i>]	Show information about all VLAN filters or about a specified VLAN or VLAN access map.



Configuring QoS

This chapter describes how to use different methods to configure quality of service (QoS) on the Catalyst 3750 Metro switch. With QoS, you can provide preferential treatment to certain traffic at the expense of others. Without QoS, the switch 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.

You can use auto-QoS to identify ports connected to Cisco IP phones and ports that receive trusted voice over IP (VoIP) traffic.

You can use standard QoS to classify, police, mark, queue, and schedule inbound traffic on any port as well as queue and schedule outbound traffic. On ingress, standard QoS offers classification based on the class of service (CoS), Differentiated Services Code Point (DSCP), or IP precedence value in the inbound packet. You can perform the classification based on Layer 2 MAC, IP-standard, or IP-extended access control lists (ACLs). Standard QoS also offers single-rate or aggregate traffic policing on ingress. Drop policy actions are passing through the packet without modification, marking down the assigned DSCP in the packet, or dropping the packet. Standard QoS performs ingress queueing based on the weighted tail drop (WTD) algorithm and ingress scheduling based on shaped round robin (SRR). On egress, standard QoS offers queueing based on WTD and scheduling based on SRR shared or shaped weights.

You can use hierarchical QoS to classify, police, mark, queue, and schedule outbound traffic on an enhanced-services (ES) port. On egress, hierarchical QoS offers classification based on the CoS, DSCP, IP precedence, or the multiprotocol label switching (MPLS) experimental (EXP) bits in the outbound packet. You also can classify a packet based on its VLAN. Hierarchical QoS offers two-rate traffic policing on egress. Drop policy actions are passing through the packet without modification; marking down the CoS, DSCP, IP precedence, or the MPLS EXP bits in the packet; or dropping the packet. Hierarchical QoS performs egress queueing based on tail drop or Weighted Random Early Detection (WRED). The queue scheduling management features is class-based weighted fair queueing (CBWFQ), and the scheduling congestion-management features is low-latency queueing (LLQ). You can use traffic shaping to decrease the burstiness of traffic.

For information about multiprotocol label switching (MPLS), Ethernet over MPLS (EoMPLS), and QoS, see the [“Configuring MPLS and EoMPLS QoS” section on page 30-18](#).



Note

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

This chapter consists of these sections:

- [Understanding QoS, page 26-2](#)
- [Understanding Hierarchical QoS, page 26-19](#)
- [Configuring Auto-QoS, page 26-29](#)
- [Displaying Auto-QoS Information, page 26-37](#)
- [Configuring Standard QoS, page 26-37](#)
- [Displaying Standard QoS Information, page 26-75](#)
- [Configuring Hierarchical QoS, page 26-76](#)
- [Displaying Hierarchical QoS Information, page 26-101](#)

Understanding QoS

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 provide preferential treatment. Implementing QoS in your network makes network performance more predictable and bandwidth utilization more effective.

The QoS implementation is based on the Differentiated Services (Diff-Serv) 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 also can be carried in the Layer 2 frame. These special bits in the Layer 2 frame or in the Layer 3 packet are described here and shown in [Figure 26-1](#):

- Prioritization bits in Layer 2 frames:

Layer 2 Inter-Switch Link (ISL) frame headers have a 1-byte User field that carries an IEEE 802.1p CoS value in the three least-significant bits. On ports configured as Layer 2 ISL trunks, all traffic is in ISL frames.

Layer 2 802.1Q frame headers have a 2-byte Tag Control Information field that carries the CoS value in the three most-significant bits, which are called the User Priority bits. On ports 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 either an IP precedence value or a DSCP value. QoS supports the use of either value because DSCP values are backward-compatible with IP precedence values.

IP precedence values range from 0 to 7.


DSCP values range from 0 to 63.

Figure 26-1 QoS Classification Layers in Frames and Packets


Encapsulated Packet




Layer 2 ISL Frame



Layer 2 802.1Q and 802.1p Frame



Layer 3 IPv4 Packet

**Note**

Layer 3 IPv6 packets are treated as non-IP packets and are bridged by the switch.

All switches and routers that access the Internet rely on the class information to provide 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 with this task.

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 provide 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 inbound and outbound traffic.

Basic QoS Model

To implement QoS, the switch must distinguish (classify) packets or flows from one another, assign a label to indicate the given quality of service as the packets move through the switch, make the packets comply with the configured resource usage limits (police and mark), and provide different treatment (queue and schedule) in all situations where resource contention exists. The switch also needs to ensure that traffic sent from it meets a specific traffic profile (shape).

Figure 26-2 shows the basic QoS model.

Figure 26-2 Basic QoS Model



98887

These are the actions when traffic is received by the switch:

- Classification is the process of generating a distinct path for a packet by associating it with a QoS label. The switch maps the CoS or DSCP in the packet to a QoS label to distinguish one kind of traffic from another. The QoS label that is generated identifies all future QoS actions to be performed on this packet. For more information, see the [“Ingress Classification” section on page 26-6](#).
- Policing decides whether a packet is in or out of profile by comparing the rate of the inbound traffic to the configured policer. The policer limits the bandwidth consumed by a flow of traffic. The result is passed to the marker. For more information, see the [“Ingress Policing and Marking” section on page 26-9](#).

- Marking evaluates the policer configuration information for the action to take when a packet is out of profile. Marking actions are to pass through a packet without modification, to mark down the QoS label in the packet, or to drop the packet. For more information, see the [“Ingress Policing and Marking” section on page 26-9](#).
- Queueing evaluates the QoS label and the corresponding DSCP or CoS value to select into which of the two ingress queues to place a packet. Queueing is enhanced with the WTD algorithm, a congestion-avoidance mechanism. If the threshold is exceeded, the packet is dropped. For more information, see the [“Queueing and Scheduling Overview” section on page 26-12](#).
- Scheduling services the queues based on their configured (SRR) weights. One of the ingress queues is the priority queue, and SRR services it for its configured share before servicing the other queue. For more information, see the [“SRR Shaping and Sharing” section on page 26-14](#).

These are the actions when the traffic is sent out a standard port:

- Queueing evaluates the QoS label and the corresponding DSCP or CoS value to select into which queue-set (a set of four queues per port) to place a packet. Because congestion can occur when multiple ingress ports simultaneously send data to an egress port, WTD is used to differentiate traffic classes and to subject the packets to different thresholds based on the QoS label. If the threshold is exceeded, the packet is dropped. For more information, see the [“Queueing and Scheduling Overview” section on page 26-12](#).
- Scheduling services the four egress queues based on their configured SRR shared or shaped weights.

These are the actions when the traffic is sent out an ES port:

- Classification is the process of generating a distinct path for a packet by matching the CoS, DSCP, IP precedence, MPLS EXP bits in the header, or matching a packet based on the inner and the outer VLAN IDs. The hierarchical configuration controls the number of class-level, VLAN-level, and physical-interface-level queues to create. For information, see the [“Understanding Hierarchical QoS” section on page 26-19](#) and the [“Hierarchical Levels” section on page 26-20](#). For classification information, see the [“Egress Classification Based on Traffic Classes and Traffic Policies” section on page 26-23](#).
- Policing decides whether a packet is in or out of profile by comparing the rate of the outbound traffic to the configured policer. The policer limits the bandwidth consumed by a flow of traffic. The result is passed to the marker. For more information, see the [“Egress Policing and Marking” section on page 26-24](#).
- 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 QoS label in the packet, or drop the packet). For more information, see the [“Egress Policing and Marking” section on page 26-24](#).
- Queueing is accomplished through a hierarchical queueing framework, in which the switch assigns each packet to a queue based on the packet class, VLAN, and physical interface. Tail drop or WRED can be configured per queue as the congestion-avoidance mechanism. With tail drop, packets are queued until the maximum threshold is exceeded, and then all the packets are dropped. WRED reduces the chances of tail drop by selectively dropping packets when the port begins to show signs of congestion. For more information, see the [“Queueing and Scheduling of Hierarchical Queues” section on page 26-26](#).
- Scheduling is accomplished through CBWFQ or LLQ (strict priority queueing). CBWFQ is a mechanism that provides guaranteed bandwidth to a particular traffic class while still fairly serving all other traffic in the network. LLQ is another scheduling mechanism, which ensures that delay-sensitive traffic is queued and sent before the traffic in other queues. Scheduling services the queues through average-rate shaping. For more information, see the [“Queueing and Scheduling of Hierarchical Queues” section on page 26-26](#).

Ingress Classification

Ingress classification is the process of distinguishing one kind of traffic from another by examining the fields in the packet upon receipt. Classification is enabled only if QoS is globally enabled on the switch. By default, QoS is globally disabled, so no classification occurs.

During ingress classification, the switch performs a lookup and assigns a QoS label to the packet. The QoS label identifies all QoS actions to be performed on the packet and from which queue the packet is sent.

The QoS label is based on the DSCP or the CoS value in the packet and decides the queueing and scheduling actions to perform on the packet. The label is mapped according to the trust setting and the packet type as shown in [Figure 26-3 on page 26-7](#).

You specify which fields in the frame or packet that you want to use to classify inbound traffic. For non-IP traffic, you have these ingress classification options as shown in [Figure 26-3](#):

- Trust the CoS value in the inbound frame (configure the port to trust CoS). Then use the configurable CoS-to-DSCP map to generate a DSCP value for the packet. Layer 2 ISL frame headers carry the CoS value in the three least-significant bits of the 1-byte User field. 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.
- Trust the DSCP or trust IP precedence value in the inbound frame. These configurations are meaningless for non-IP traffic. If you configure a port with either of these options and non-IP traffic is received, the switch assigns a CoS value and generates a DSCP value from the CoS-to-DSCP map.
- Perform the classification based on a configured Layer 2 MAC ACL, which can examine the MAC source address, the MAC destination address, and other fields. If no ACL is configured, the packet is assigned 0 as the DSCP and CoS values, which means best-effort traffic. Otherwise, the policy-map action specifies a DSCP or a CoS value to assign to the inbound frame.

For IP traffic, you have these ingress classification options as shown in [Figure 26-3](#):

- Trust the DSCP value in the inbound packet (configure the port to trust DSCP), and assign the same DSCP value to the packet. The IETF defines the six most-significant bits of the 1-byte ToS field as the DSCP. The priority represented by a particular DSCP value is configurable. DSCP values range from 0 to 63.

For ports that are on the boundary between two QoS administrative domains, you can modify the DSCP to another value through the configurable DSCP-to-DSCP-mutation map.


- Trust the IP precedence value in the inbound packet (configure the port to trust IP precedence), and generate a DSCP value for the packet through the configurable IP-precedence-to-DSCP map. The IP version 4 specification defines the three most-significant bits of the 1-byte ToS field as the IP precedence. IP precedence values range from 0 for low priority to 7 for high priority.
- Trust the CoS value (if present) in the inbound packet, and generate a DSCP value for the packet through the CoS-to-DSCP map. If the CoS value is not present, use the default port CoS value.
- Perform the classification based on a configured IP standard or an extended ACL, which examines various fields in the IP header. If no ACL is configured, the packet is assigned 0 as the DSCP and CoS values, which means best-effort traffic. Otherwise, the policy-map action specifies a DSCP or a CoS value to assign to the inbound frame.

For information on the maps described in this section, see the [“Mapping Tables” section on page 26-11](#). For configuration information on port trust states, see the [“Configuring Ingress Classification by Using Port Trust States” section on page 26-42](#).

You can configure classification through ACLs, traffic classes, and traffic policies. For more information, see the [“Ingress Classification Based on QoS ACLs”](#) section on page 26-8 and the [“Ingress Classification Based on Traffic Classes and Traffic Policies”](#) section on page 26-8.

After ingress classification, the packet is sent to the policing, marking, and the ingress queueing and scheduling stages.

Figure 26-3 Ingress Classification Flowchart



86834

Ingress Classification Based on QoS ACLs

You can use IP standard, IP extended, or Layer 2 MAC 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 a match with a deny action is encountered, the ACL being processed is skipped, and the next ACL is processed.
- If no match with a permit action is encountered and all the ACEs have been examined, no QoS processing occurs on the packet, and the switch offers best-effort service to the packet.
- If multiple ACLs are configured on a port, the lookup stops after the packet matches the first ACL with a permit action, and QoS processing begins.

**Note**

When you create 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.

After a traffic class is 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 to rate-limit the class. This ingress policy is then attached to a port.

You implement IP ACLs to classify IP traffic by using the **access-list** global configuration command; you implement Layer 2 MAC ACLs to classify non-IP traffic by using the **mac access-list extended** global configuration command. For configuration information, see the [“Configuring an Ingress QoS Policy” section on page 26-48](#).

Ingress Classification Based on Traffic Classes and Traffic Policies

You define a traffic class to classify traffic, use a traffic policy to decide how to treat the classified traffic, and attach the ingress policy to a port to create a service policy.

You use the class map to define a specific traffic flow (or class) and to isolate it 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 or matching a specific list of DSCP or IP precedence values. 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.

You create a class map by using the **class-map** global configuration command. When you enter the **class-map** 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. Inbound packets are checked against the match criteria configured for a class map to decide if the packet belongs to that class. If a packet matches the specified criteria, the packet is considered a member of the class and is forwarded according to the QoS specifications set in the traffic policy. If a packet fails to meet any of the matching criteria, it is classified as a member of the default traffic class if one is configured.

You use the policy map to create the traffic policy, to specify the traffic class to act on, and to configure the QoS features associated with the traffic class. Actions on ingress can include trusting the received CoS, DSCP, or IP precedence values in the traffic class; setting a specific DSCP or IP precedence value in the traffic class; or specifying the traffic bandwidth limitations and the action to take when the traffic is out of profile.

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 use the **class** policy-map configuration command to name the traffic class associated with the traffic policy. If you specify **class-default** as the class name in the **class** policy-map configuration command, packets that fail to meet any of the matching criteria are classified as members of the default traffic class. You can manipulate this class (for example, police it and mark it) just like any traffic class, but you cannot delete it. After you name the traffic class with the **class** command, the switch enters policy-map class configuration mode, and you can specify the actions to take on this traffic class.

An ingress policy-map can include **police**, **police aggregate**, **trust**, or **set** policy-map class configuration commands. You attach the ingress policy-map to a port by using the **service-policy input** interface configuration command.

For more information, see the “[Ingress Policing and Marking](#)” section on page 26-9. For configuration information, see the “[Configuring an Ingress QoS Policy](#)” section on page 26-48.

Ingress Policing and Marking

After a packet is classified and has a DSCP-based or CoS-based QoS label assigned to it, the traffic policing and marking process can begin. [Figure 26-4 on page 26-11](#) shows an ingress, single-rate traffic policer.

Ingress traffic policing controls the maximum rate of traffic received on a port. The policer defines the bandwidth limitations of the traffic, and the action to take if the limits are exceeded. It is often configured on ports at the edge of a network to limit traffic into or out of the network. In most policing configurations, traffic that falls within the rate parameters is sent. Traffic that exceeds the parameters is considered to be *out of profile* or *nonconforming* and is dropped or sent with a different priority.



Note

All traffic, regardless of whether it is bridged or routed, is subjected to a configured policer. As a result, bridged packets might be dropped or might have their DSCP or CoS fields modified when they are policed and marked.

You can create these types of ingress policers:

- Individual (single-rate)

QoS applies the bandwidth limits specified in the policer separately to each matched traffic class. You can configure a single-rate policer within a policy map by using the **police** policy-map class configuration command.

- Aggregate

QoS applies the bandwidth limits specified in an aggregate policer cumulatively to all matched traffic flows. You configure this type of policer by specifying the aggregate policer name within a policy map by using the **police aggregate** policy-map class configuration command. You specify the bandwidth limits of the policer by using the **mls qos aggregate-policer** global configuration command. In this way, the aggregate policer is shared by multiple classes of traffic within a policy map.


A single-rate traffic policer decides on a packet-by-packet basis whether the packet is in or out of profile and specifies the actions on the packet. These actions, carried out by the marker, include passing through the packet without modification, dropping the packet, or modifying (marking down) the assigned DSCP of the packet and allowing the packet to pass through. The configurable policed-DSCP map provides the packet with a new DSCP-based QoS label. For information on the policed-DSCP map, see the [“Mapping Tables” section on page 26-11](#). Marked-down packets use the same queues as the original QoS label to prevent packets in a flow from getting out of order.

Single-rate policing uses a token-bucket algorithm. As each frame is received by the switch, a token is added to the bucket. The bucket has a hole in it and leaks at a rate that you specify as the average traffic rate in bps. Each time a token is added to the bucket, the switch verifies that there is enough room in the bucket. If there is not enough room, the packet is marked as nonconforming, and the specified policer action is taken (dropped or marked down).

How quickly the bucket fills is a function of the bucket depth (burst-byte), the rate at which the tokens are removed (rate-bps), and the duration of the burst above the average rate. The size of the bucket imposes an upper limit on the burst length and limits the number of frames that can be sent back-to-back. If the burst is short, the bucket does not overflow, and no action is taken against the traffic flow. However, if a burst is long and at a higher rate, the bucket overflows, and the policing actions are taken against the frames in that burst.

You configure the bucket depth (the maximum burst that is tolerated before the bucket overflows) by using the *burst-byte* option of the **police** policy-map class configuration command or the **mls qos aggregate-policer** global configuration command. You configure how fast (the average rate) that the tokens are removed from the bucket by using the *rate-bps* option of the **police** policy-map class configuration command or the **mls qos aggregate-policer** global configuration command.

After you configure the policy map and policing actions, attach the ingress policy to a port by using the **service-policy input** interface configuration command. For configuration information, see the [“Classifying, Policing, and Marking Ingress Traffic by Using Policy Maps” section on page 26-54](#) and the [“Classifying, Policing, and Marking Ingress Traffic by Using Aggregate Policers” section on page 26-57](#).

Figure 26-4 Ingress, Single-Rate Policing and Marking Flowchart

86835

Mapping Tables

During QoS ingress processing, the switch represents the priority of all traffic (including non-IP traffic) with an QoS label based on the DSCP or CoS value from the classification stage:

- During ingress classification, QoS uses configurable mapping tables to derive a corresponding DSCP or CoS value from a received CoS, DSCP, or IP precedence value. These maps include the CoS-to-DSCP map and the IP-precedence-to-DSCP map. You configure these maps by using the **mls qos map cos-dscp** and the **mls qos map ip-prec-dscp** global configuration commands.

On an ingress port configured in the DSCP-trusted state, if the DSCP values are different between the QoS domains, you can apply the configurable DSCP-to-DSCP-mutation map to the port that is on the boundary between the two QoS domains. You configure this map by using the **mls qos map dscp-mutation** global configuration command.

- During ingress policing, QoS can assign another DSCP value to an IP or a non-IP packet (if the packet is out of profile and the policer specifies a marked-down value). This configurable map is called the policed-DSCP map. You configure this map by using the **mls qos map policed-dscp** global configuration command.

- Before the traffic reaches the scheduling stage, QoS stores the packet in an ingress and an egress queue according to the QoS label. The QoS label is based on the DSCP or the CoS value in the inbound packet. The QoS label selects an ingress queue and an egress queue from the queue-set through the DSCP input and output queue threshold maps or through the CoS input and output queue threshold maps. You configure these maps by using the **mls qos srr-queue {input | output} dscp-map** and the **mls qos srr-queue {input | output} cos-map** global configuration commands.

The CoS-to-DSCP, DSCP-to-CoS, and the IP-precedence-to-DSCP maps have default values that might or might not be appropriate for your network.

The default DSCP-to-DSCP-mutation map and the default policed-DSCP map are null maps; they map an inbound DSCP value to the same DSCP value. The DSCP-to-DSCP-mutation map is the only map you apply to a specific port. All other maps apply to the entire switch.


For configuration information, see the [“Configuring DSCP Maps” section on page 26-59](#).

For information about the DSCP and CoS input queue threshold maps, see the [“Queueing and Scheduling of Ingress Queues” section on page 26-15](#). For information about the DSCP and CoS output queue threshold maps, see the [“Queueing and Scheduling of Egress Queue-Sets” section on page 26-17](#).

Queueing and Scheduling Overview

The switch has queues at specific points to help prevent congestion, as shown in [Figure 26-5](#).

Figure 26-5 Ingress and Egress Queue Location



Because the total ingress bandwidth of all ports can exceed the bandwidth of the internal ring, ingress queues are located after the packet is classified, policed, and marked and before packets are forwarded into the switch fabric. Because multiple ingress ports can simultaneously send packets to an egress port and cause congestion, egress queue-sets are located after the internal ring.

Each port belongs to an egress queue-set, which defines all the characteristics of the four queues per port. The ES ports also use a hierarchical queueing model in which each packet is assigned to a hierarchical queue based on the physical interface, VLAN, or class. Traffic destined for an ES port passes through the queue-set before reaching the hierarchical queues. If congestion occurs in the hierarchical queues and backs up to the queue-sets, the queue-set configuration controls how traffic is dropped.

Ingress queues use WTD for congestion management. The egress queue-sets also use WTD. For more information, see the next section.

Ingress queues support SRR in shared mode for scheduling. The egress queue-sets also support SRR in shared or shaped mode for scheduling. For more information, see the [“SRR Shaping and Sharing” section on page 26-14](#).

For information about ingress queueing and scheduling, see the [“Queueing and Scheduling of Ingress Queues” section on page 26-15](#). For information about egress queue-set queueing and scheduling, see the [“Queueing and Scheduling of Egress Queue-Sets” section on page 26-17](#).

The egress hierarchical queues for ES ports use tail drop or WRED for congestion management. These ports also use CBWFQ or LLQ for scheduling. For more information, see the [“Queueing and Scheduling of Hierarchical Queues” section on page 26-26](#).

Weighted Tail Drop


The ingress queues and the egress queue-sets use an enhanced version of the tail-drop congestion-avoidance mechanism called WTD. WTD manages the queue lengths and provides drop precedences for different traffic classifications.

As a frame is sent to a particular queue, WTD uses the frame’s assigned QoS label to subject it to different thresholds. If the threshold is exceeded for that QoS label (the space available in the destination queue is less than the size of the frame), the switch drops the frame.

[Figure 26-6](#) shows an example of WTD operating on a queue whose size is 1000 frames. Three drop percentages are configured: 40, 60, and 100 percent. These percentages mean that up to 400 frames can be queued at the 40-percent threshold, up to 600 frames at the 60-percent threshold, and up to 1000 frames at the 100-percent threshold.

In this example, CoS values 6 and 7 have a greater importance than the other CoS values, and they are assigned to the 100-percent drop threshold (queue-full state). CoS values 4 and 5 are assigned to the 60-percent threshold, and CoS values 0 to 3 are assigned to the 40-percent threshold.

Suppose the queue is already filled with 600 frames, and a new frame arrives. It contains CoS values 4 and 5 and is subjected to the 60-percent threshold. If this frame is added to the queue, the threshold will be exceeded, so the switch drops it.

Figure 26-6 WTD and Queue Operation

For more information, see the [“Mapping DSCP or CoS Values to an Ingress Queue and Setting WTD Thresholds”](#) section on page 26-65, the [“Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set”](#) section on page 26-69, and the [“Mapping DSCP or CoS Values to an Egress Queue-Set and to a Threshold ID”](#) section on page 26-71.

SRR Shaping and Sharing

The ingress queues and egress queue-sets are serviced by SRR, which controls the rate at which packets are sent. On the ingress queues, SRR sends packets to the internal ring. On the egress queue-sets, SRR sends packets to a standard port.

You can configure SRR on the egress queue-sets for sharing or for shaping. However, for ingress queues, sharing is the default mode and is the only mode supported.

In shaped mode, the queues are guaranteed a percentage of the bandwidth, and they are rate-limited to that amount. Shaped traffic does not use more than the allocated bandwidth even if the link is idle. Shaping provides a more even flow of traffic over time and reduces the peaks and valleys of bursty traffic. With shaping, the absolute value of each weight is used to compute the bandwidth available for the queues.


In shared mode, the queues share the bandwidth among them according to the configured weights. The bandwidth is guaranteed at this level but not limited to it. For example, if a queue is empty and no longer requires a share of the link, the remaining queues can expand into the unused bandwidth and share it among them. With sharing, the ratio of the weights controls the frequency of dequeuing; the absolute values are meaningless.

For more information, see the [“Allocating Bandwidth Between the Ingress Queues”](#) section on page 26-67, the [“Configuring SRR Shaped Weights on an Egress Queue-Set”](#) section on page 26-72, and the [“Configuring SRR Shared Weights on an Egress Queue-Set”](#) section on page 26-73.

Queueing and Scheduling of Ingress Queues

Figure 26-7 shows the ingress queueing and scheduling flowchart.

Figure 26-7 Ingress Queueing and Scheduling Flowchart



Note

SRR services the priority queue for its configured share before servicing the other queue.

The switch supports two configurable ingress queues, which are serviced by SRR in shared mode only. Table 26-1 describes the queues.

Table 26-1 Ingress Queue Types

Queue Type ¹	Function
Normal	User traffic that is considered to be normal priority. You can configure three different thresholds to differentiate among the flows. You can use the mls qos srr-queue input threshold , the mls qos srr-queue input dscp-map , and the mls qos srr-queue input cos-map global configuration commands.
Expedite	High-priority user traffic such as differentiated services (DF) expedited forwarding or voice traffic. You can configure the bandwidth required for this traffic as a percentage of the total traffic by using the mls qos srr-queue input priority-queue global configuration command. The expedite queue has guaranteed bandwidth.

1. The switch uses two nonconfigurable queues for traffic that is essential for proper network operation.

You assign each packet that flows through the switch to a queue and to a threshold. Specifically, you map DSCP or CoS values to an ingress queue and map DSCP or CoS values to a threshold ID. You use the **mls qos srr-queue input dscp-map queue** *queue-id* {*dscp1...dscp8* | **threshold** *threshold-id* *dscp1...dscp8*} or the **mls qos srr-queue input cos-map queue** *queue-id* {*cos1...cos8* | **threshold** *threshold-id* *cos1...cos8*} global configuration command. You can display the DSCP input queue threshold map and the CoS input queue threshold map by using the **show mls qos maps** privileged EXEC command.

WTD Thresholds

The queues use WTD to support distinct drop percentages for different traffic classes. Each queue has three drop thresholds: two configurable (*explicit*) WTD thresholds and one nonconfigurable (*implicit*) threshold preset to the queue-full state. You assign the two explicit WTD threshold percentages for threshold ID 1 and ID 2 to the ingress queues by using the **mls qos srr-queue input threshold** *queue-id* *threshold-percentage1* *threshold-percentage2* global configuration command. Each threshold value is a percentage of the total number of allocated buffers for the queue. The drop threshold for threshold ID 3 is preset to the queue-full state, and you cannot modify it. For more information about how WTD works, see the [“Weighted Tail Drop” section on page 26-13](#).

Buffer and Bandwidth Allocation

You define the ratio (allocate the amount of space) with which to divide the ingress buffers between the two queues by using the **mls qos srr-queue input buffers** *percentage1* *percentage2* global configuration command. The buffer allocation together with the bandwidth allocation control how much data can be buffered and sent before packets are dropped. You allocate bandwidth as a percentage by using the **mls qos srr-queue input bandwidth** *weight1* *weight2* global configuration command. The ratio of the weights is the ratio of the frequency in which the SRR scheduler sends packets from each queue to the internal ring.

Priority Queueing

You can configure one ingress queue as the priority queue by using the **mls qos srr-queue input priority-queue** *queue-id* **bandwidth** *weight* global configuration command. The priority queue should be used for traffic (such as voice) that requires guaranteed delivery because this queue is guaranteed part of the bandwidth regardless of the load on the internal ring.


SRR services the priority queue for its configured weight as specified by the **bandwidth** keyword in the **mls qos srr-queue input priority-queue** *queue-id* **bandwidth** *weight* global configuration command. Then, SRR shares the remaining bandwidth with both ingress queues and services them as specified by the weights configured with the **mls qos srr-queue input bandwidth** *weight1* *weight2* global configuration command.

You can combine the commands described in this section to prioritize traffic by placing packets with particular DSCP or CoS values into certain queues, by allocating a large queue size or by servicing the queue more frequently, and by adjusting queue thresholds so that packets with lower priorities are dropped. For configuration information, see the [“Configuring Ingress Queue Characteristics” section on page 26-64](#).

Queueing and Scheduling of Egress Queue-Sets

Figure 26-8 shows the egress queue-set queueing and scheduling flowchart.

Figure 26-8 Egress Queue-Set Queueing and Scheduling Flowchart




98884

Each port supports four egress queues. These queues are assigned to a queue-set. All traffic exiting the switch on a standard port flows through one of these four queues and is subjected to a threshold based on the QoS label assigned to the packet. Traffic destined for an ES port passes through the queue-set before reaching the hierarchical queues. If congestion occurs in the hierarchical queues that backs up to the queue-sets, the queue-set configuration controls how traffic is dropped.

Figure 26-9 shows the egress queue-set buffer. The buffer space is divided between the common pool and the reserved pool. The switch uses a buffer allocation scheme to reserve a minimum amount of buffers for each egress queue, to prevent any queue or port from consuming all the buffers and depriving other queues, and to control whether to grant buffer space to a requesting queue. The switch detects whether or not the target queue has consumed more buffers than its reserved amount (under-limit), whether it has consumed all of its maximum buffers (over limit), and whether the common pool is empty (no free buffers) or not empty (free buffers). If the queue is not over-limit, the switch can allocate buffer space from the reserved pool or from the common pool (if it is not empty). If there are no free buffers in the common pool or if the queue is over-limit, the switch drops the frame.

Figure 26-9 Egress Queue-Set Buffer Allocation



Buffer and Memory Allocation

You guarantee the availability of buffers, set drop thresholds, and configure the maximum memory allocation for a queue-set by using the **mls qos queue-set output *qset-id* threshold *queue-id* drop-threshold1 drop-threshold2 reserved-threshold maximum-threshold** global configuration command. Each threshold value is a percentage of the queue's allocated memory, which you specify by using the **mls qos queue-set output *qset-id* buffers allocation1 ... allocation4** global configuration command. The sum of all the allocated buffers represents the reserved pool, and the remaining buffers are part of the common pool.

Through buffer allocation, you can ensure that high-priority traffic is buffered. For example, if the buffer space is 400, you can allocate 70 percent of it to queue 1 and 10 percent to queues 2 through 4. Queue 1 then has 280 buffers allocated to it, and queues 2 through 4 each have 40 buffers allocated to them.

You can guarantee that the allocated buffers are reserved for a specific queue in a queue-set. For example, if there are 100 buffers for a queue, you can reserve 50 percent (50 buffers). The switch returns the remaining 50 buffers to the common pool. You also can enable a queue in the full condition to obtain more buffers than are reserved for it by setting a maximum threshold. The switch can allocate the needed buffers from the common pool if the common pool is not empty.

WTD Thresholds

You can assign each packet that flows through the switch to a queue and to a threshold. Specifically, you map DSCP or CoS values to an egress queue-set and map DSCP or CoS values to a threshold ID. You use the **mls qos srr-queue output dscp-map queue queue-id {dscp1...dscp8 | threshold threshold-id dscp1...dscp8}** or the **mls qos srr-queue output cos-map queue queue-id {cos1...cos8 | threshold threshold-id cos1...cos8}** global configuration command. You can display the DSCP output queue threshold map and the CoS output queue threshold map by using the **show mls qos maps** privileged EXEC command.

The queue-set uses WTD to support distinct drop percentages for different traffic classes. Each queue has three drop thresholds: two configurable (*explicit*) WTD thresholds and one nonconfigurable (*implicit*) threshold preset to the queue-full state. You assign the two WTD threshold percentages for threshold ID 1 and ID 2. The drop threshold for threshold ID 3 is preset to the queue-full state, and you cannot modify it. For more information about how WTD works, see the [“Weighted Tail Drop” section on page 26-13](#).

Shaped or Shared Mode

SRR services each queue-set in shared or shaped mode. You map a port to a queue-set by using the **queue-set qset-id** interface configuration command. You assign shared or shaped weights to a standard port by using the **srr-queue bandwidth share weight1 weight2 weight3 weight4** or the **srr-queue bandwidth shape weight1 weight2 weight3 weight4** interface configuration command. You can assign only shared weights to an ES port. For an explanation of the differences between shaping and sharing, see the [“SRR Shaping and Sharing” section on page 26-14](#).

The buffer allocation together with the SRR weight ratios control how much data can be buffered and sent before packets are dropped. The weight ratio is the ratio of the frequency in which the SRR scheduler sends packets from each queue.

You can combine the commands described in this section to prioritize traffic by placing packets with particular DSCP or CoS values into certain queues, by allocating a large queue size or by servicing the queue more frequently, and by adjusting queue thresholds so that packets with lower priorities are dropped. For configuration information, see the [“Configuring Egress Queue-Set Characteristics” section on page 26-69](#).



Note

The egress queue-set default settings are suitable for most situations. You should change them only when you have a thorough understanding of the queues and only if these settings do not meet your QoS solution.

Understanding Hierarchical QoS

The switch supports a hierarchical QoS configuration (traffic classification, CBWFQ, LLQ, shaping, and two-rate three-color policing) that is applied to the output of the ES ports.

Hierarchical QoS configuration is based on the concept of a bandwidth-limited stream of traffic, which is a stream of packets that has its departure rate constrained in some manner. At each level of the hierarchy, the switch must classify each packet to select into which traffic stream in that level the packet belongs. When the stream is classified, if its arrival rate exceeds its departure rate, queueing can become congested. To compensate for this, you can configure policies that contain policer drops, configure tail drop or WRED, a congestion-avoidance technique, or to influence whether the packet is queued. You also can implement scheduling policies (CBWFQ, LLQ, and shaping) to influence how quickly a packet is sent out the port.

Hierarchical Levels

Hierarchical QoS configuration involves traffic classification, policing, queueing, and scheduling. You can create a hierarchy by associating a class-level policy-map with a VLAN-level policy-map, by associating that VLAN-level policy-map with a physical-level policy-map, and by attaching the physical-level policy-map to the output of an ES port. You can skip hierarchical levels, but the order of the levels (class level, VLAN level, and then the physical level) must be preserved.

You can configure these three QoS levels in the hierarchy:

- **Class level**—You configure this level of the hierarchy by matching CoS, DSCP, IP precedence, or MPLS EXP bits in the outbound packet through the **match {cos cos-list | ip dscp dscp-list | ip precedence ip-precedence-list | mpls experimental exp-list}** class-map configuration command. At the class level, you can:
 - Configure policer drops by using the **police cir** or **police cir percent** policy-map class configuration command.
 - Configure tail drop or WRED drop policies by using the **queue-limit** or the **random-detect** policy-map class configuration command.
 - Modify the traffic class by setting Layer 2 and Layer 3 QoS fields through the **set {cos new-cos | ip {dscp new-dscp | precedence new-precedence} | mpls experimental exp-number}** policy-map class configuration command.
 - Configure CBWFQ or LLQ scheduling by using the **bandwidth** or the **priority** policy-map class configuration command.
 - Configure egress traffic shaping by using the **shape** policy-map class configuration command.

The switch supports eight classes (including the default class) per policy map at this level. The default class is reserved for packets that do not meet any of the matching criteria.

This is an example of a class-level classification and its naming convention:

```
Switch(config)# class-map match-any class-level-class-map-name
Switch(config-cmap)# match ip dscp 10 11 12
```

This is an example of a class-level policy-map and its naming convention:

```
Switch(config)# policy-map class-level-policy-map-name
Switch(config-pmap)# class class-level-class-name
Switch(config-pmap-c)# bandwidth percent 20
Switch(config-pmap-c)# shape average 20000000
```

This is a class-level configuration example that combines a class-level classification and a class-level policy-map to create a service policy:

```
Switch(config)# class-map c1
Switch(config-cmap)# match ip precedence 4
Switch(config-cmap)# exit
Switch(config)# policy-map policy1
Switch(config-pmap)# class c1
Switch(config-pmap-c)# police cir 500000 bc 10000 pir 1000000 be 10000 conform-action
transmit exceed-action set-prec-transmit 2 violate-action drop
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/1
Switch(config-if)# service-policy output policy1
```


- VLAN level—You start configuration of per-VLAN QoS by entering the **match vlan** *vlan-id* class-map configuration command on one or more VLANs or by entering the **match vlan** *vlan-id* and the **match vlan inner** *vlan-id* class-map configuration command on one or more 802.1Q tunnels. At this level, you can configure the VLANs or 802.1Q tunnels to police, to share the available port bandwidth and to enable CBWFQ, and to shape the traffic. You configure these features by using the **police cir**, **police cir percent**, **bandwidth**, and **shape** policy-map class configuration commands.

For a finer level of control, you also can associate a previously defined child policy at the class level with a new service policy by using the **service-policy** policy-map class configuration command. In the class-level child policy, you can configure tail drop or WRED drop policies, set Layer 2 and Layer 3 QoS fields, or enable the priority queue. These features are available only at the class level. By using a child policy, you apply a class-level policy only to traffic that matches the VLAN class.

You cannot mix VLAN-level and class-level matches within a class map.

You can attach up to 2045 user-created VLAN-level classes. This means that you can have 1022 unique classes and can associate them with the two ES ports (and have one left over), or you can add more classes to one ES port and can subtract from the other one. You can shape every class that you configure. You can create up to 4093 class maps.

This is an example of a VLAN-level classification and its naming convention:

```
Switch(config)# class-map match-all vlan-level-class-map-name
Switch(config-cmap)# match vlan 5
Switch(config-cmap)# match vlan inner 3 - 8
```

This is an example of a VLAN-level policy-map and its naming convention:

```
Switch(config)# policy-map vlan-level-policy-map-name
Switch(config-pmap)# class vlan-level-class-name
Switch(config-pmap-c)# police cir 500000 bc 10000 pir 1000000 be 10000 conform-action
transmit exceed-action set-prec-transmit 2 violate-action drop
```

This is an example of a VLAN-level policy-map and its naming convention when a previously defined child policy is associated at the class level:

```
Switch(config)# policy-map vlan-level-policy-map-name
Switch(config-pmap)# class vlan-level-class-name
Switch(config-pmap-c)# bandwidth percent 30
Switch(config-pmap-c)# service-policy class-level-policy-map-name
```

This is a VLAN-level configuration example that combines a VLAN-level classification and a VLAN-level policy-map:

```
Switch(config)# class-map match-all vlan203
Switch(config-cmap)# match vlan 203
Switch(config-cmap)# exit
Switch(config)# policy-map vlan-policy
Switch(config-pmap)# class vlan203
Switch(config-pmap-c)# police cir 500000 bc 10000 pir 1000000 be 10000 conform-action
transmit exceed-action set-prec-transmit 2 violate-action drop
```

This is an example of a VLAN-level policy-map that combines a VLAN-level classification with a VLAN-level policy-map and associates a previously defined child policy at the class level:

```
Switch(config)# class-map cls-class
Switch(config-cmap)# match mpls experimental 2
Switch(config-cmap)# exit
Switch(config)# class-map log-class
Switch(config-cmap)# match vlan 203
Switch(config-cmap)# exit
Switch(config)# policy-map cls-policy
Switch(config-pmap)# class cls-class
Switch(config-pmap-c)# set mpls experimental 5
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# policy-map log-policy
Switch(config-pmap)# class log-class
Switch(config-pmap-c)# service-policy cls-policy
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/2
Switch(config-if)# switchport trunk encapsulation isl
Switch(config-if)# switchport mode trunk
Switch(config-if)# service-policy output log-policy
```

- Physical level—You can shape or police only the class-default class at the physical level of the hierarchy by using the **shape**, **police cir**, or **police cir percent** policy-map class configuration command.

Within a policy-map, the class-default applies to all traffic that is not explicitly matched within the policy map but does match the parent policy. If no parent policy is configured, the parent policy represents the physical port. In a physical-level policy-map, class-default is the only class that you can configure.

You use the **service-policy output** interface configuration command to attach an egress policy to an ES port.

This is an example of a physical-level configuration. All hierarchical levels exist, and the order is preserved. The class level at the bottom, the VLAN level in the middle, and the physical level at the top.

```
Switch(config)# class-map my-class
Switch(config-cmap)# match ip precedence 1
Switch(config-cmap)# exit
Switch(config)# class-map my-logical-class
Switch(config-cmap)# match vlan 5
Switch(config-cmap)# exit
Switch(config)# policy-map my-class-policy
Switch(config-pmap)# class my-class
Switch(config-pmap-c)# set ip precedence 2
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# policy-map my-logical-policy
Switch(config-pmap)# class my-logical-class
Switch(config-pmap-c)# shape average 400000000
Switch(config-pmap-c)# service-policy my-class-policy
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# policy-map my-physical-policy
Switch(config-pmap)# class class-default
Switch(config-pmap-c)# shape average 500000000
Switch(config-pmap-c)# service-policy my-logical-policy
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
```

```
Switch(config)# interface gigabitethernet1/1/1
Switch(config-if)# service-policy output my-physical-policy
```

Egress Classification Based on Traffic Classes and Traffic Policies

Egress classification is the process of distinguishing one kind of traffic from another by examining the fields in the packet, processing the **match** commands, and creating the egress queues before the switch sends the packet out the ES port.

You use the class map to define a specific traffic flow (or class) and to isolate it from all other traffic. The class map defines the criteria used to match against a specific traffic flow to further classify it. At the class level, the criteria can include matching CoS, DSCP, IP precedence, or MPLS EXP bits in the header. At the VLAN level, the criteria can include matching a packet based on the inner and the outer VLAN IDs. 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.

You create a class map by using the **class-map** global configuration command. When you enter the **class-map** 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. Outbound packets are checked against the match criteria configured for a class map. If a packet matches the specified criteria, the packet is considered a member of the class, the switch creates a queue for it, and the packet is forwarded according to the QoS specifications set in the traffic policy. If a packet fails to meet any of the matching criteria, it is classified as a member of the default traffic class if one is configured.

You use the policy map to create the traffic policy, to specify the traffic class to act on, and to configure the QoS features associated with the traffic class. Actions can include trusting the received CoS, DSCP, or IP precedence bits in the traffic class; setting specific CoS, DSCP, IP precedence, or MPLS EXP bits in the traffic class; or specifying the traffic bandwidth limitations and the action to take when the traffic is out of profile.

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 use the **class** policy-map configuration command to name the traffic class associated with the traffic policy. If you specify **class-default** as the class name in the **class** policy-map configuration command, packets that fail to meet any of the matching criteria are classified as members of the default traffic class. You can manipulate this class (for example, police it and mark it) just like any traffic class, but you cannot delete it.

Within a policy-map, the class-default designates all traffic that is not explicitly matched within the policy-map but does match the policy map of the parent policy. If no parent policy is configured, the parent policy represents the physical port. In the physical-level policy-map, class-default is the only class that can be configured.

After you name the traffic class with the **class** command, the switch enters policy-map class configuration mode, and you can specify the actions to take on this traffic class.

You attach an egress policy-map to an ES port by using the **service-policy output** interface configuration command. The egress policy-map can include the **bandwidth**, **police cir**, **police cir percent**, **priority**, **queue-limit**, **random-detect**, **shape**, or **set** policy-map class configuration commands. If the policy map contains the class-default class, you can configure settings only through the **police cir**, **police cir percent**, and **shape** commands.

For more information, see the [“Egress Policing and Marking”](#) section on page 26-24 and the [“Queueing and Scheduling of Hierarchical Queues”](#) section on page 26-26.

Egress Policing and Marking

Egress traffic policing controls the maximum rate of traffic sent on a port. A policer defines the bandwidth limitations of the traffic and the action to take if the limits are exceeded. It is often configured on ports at the edge of a network to limit traffic into or out of the network. In most policing configurations, traffic that falls within the rate parameters is sent. Traffic that exceeds the parameters is considered to be *out of profile* or *nonconforming* and is dropped or sent with a different priority.


You can configure a two-rate traffic policer within a policy map at the class level, at the VLAN level, and at the physical level by using the **police cir** or the **police cir percent** policy-map class configuration command. At the physical level of the hierarchy, you can police only the class-default class in an egress policy attached to an ES port.

You can configure a two-rate traffic policer to limit the transmission rate of a traffic class and mark actions (conform, exceed, and violate) for each packet. Within the conform, exceed, and violate categories, you decide packet treatments. In the most common configurations, you configure packets that conform to be sent, packets that exceed to be sent with a decreased priority, and packets that violate to be dropped. You can decrease the priority of the CoS, the DSCP, the IP precedence, or the MPLS EXP bits in the packet.

The two-rate policer manages the maximum rate of traffic through a token-bucket algorithm. The algorithm uses the configured committed information rate (CIR) and the peak information rate (PIR) rate values to control the maximum rate of traffic allowed on a port at a given moment in time. The algorithm is affected by all traffic leaving the port, and it manages network bandwidth when several large packets are sent in the same traffic stream.

A token bucket is provided for the CIR and the PIR as shown in [Figure 26-10](#).

Figure 26-10 Egress Two-Rate Policing and Marking Flowchart



You configure the CIR and PIR rates in bps (or as a percentage of the bandwidth available on an ES port), and these rates control how fast the bucket fills (is updated) with tokens. The conform burst size (bc) and the peak burst size (be) represent the depth of the CIR and PIR buckets in bytes. This depth limits the number of tokens that the bucket can accumulate. If the bucket fills to capacity, newly arriving tokens are discarded.

Each token is permission for the source to send a certain number of bits into the network. To send a packet, the number of tokens equal to the packet size must be drained from the bucket. If there are enough tokens in the bucket, the packet conforms and can pass to the next stage. Otherwise, the exceed action associated with the bucket is applied to the packet. The packet might be dropped, or its priority value might be marked down.

In this token-bucket example, if the CIR rate is 2 kbps, 2000 tokens are added to the bucket every second (for this example, consider each token to represent a single bit of information). If a 1500-byte packet arrives, 12000 tokens (1500 bytes x 8 bits per byte) must be in the bucket for the packet to pass to the next state without triggering the exceed action. If enough tokens are in the bucket, they are drained, and the packet conforms and passes to the next stage. If there are less than 12000 tokens in the bucket, the exceed action is applied to the packet. The deeper the bucket, the more data can burst through at a rate greater than the rate at which the bucket is filling. For example, if the CIR bucket holds 6000 tokens, 750 bytes of traffic can instantaneously burst without draining the bucket (and without triggering an exceed action), even though the instantaneous burst is at a greater rate than the CIR rate of 2000 bps.

If the burst sizes approach the system maximum transmission unit (MTU), the policer strictly enforces the CIR and PIR. Normal traffic jitter can cause some percentage of inbound traffic to be flagged as nonconforming even if the average inbound rate appears to conform. If the burst size is very large, on the other hand, large traffic bursts at nonconforming data rates can be passed through the policer and flagged as conforming. Setting the burst sizes too low can result in less traffic than expected, and setting them too high can result in more traffic than expected.

For packet marking actions, if the CIR is 100 kbps, the PIR is 200 kbps, and a data stream with a rate of 250 kbps arrives at the two-rate policer:

- 100 kbps is marked as conforming to the rate.
- 100 kbps is marked as exceeding the rate.
- 50 kbps is marked as violating the rate.

If you set the CIR equal to the PIR, a traffic rate that is less than the CIR or that meets the CIR is in the conform range. Traffic that exceeds the CIR rate is in the violate range.


If you set the PIR greater than the CIR, a traffic rate less than the CIR is in the conform range. A traffic rate that exceeds the CIR but is less than or equal to the PIR is in the exceed range. A traffic rate that exceeds the PIR is in the violate range.

After you configure the policy map and policing actions, attach the egress policy to an ES port by using the **service-policy output** interface configuration command. For configuration information, see the [“Configuring Hierarchical QoS” section on page 26-76](#).

Queueing and Scheduling of Hierarchical Queues

Figure 26-11 shows the egress queueing and scheduling flowchart for a hierarchical queue.

Figure 26-11 Egress Queueing and Scheduling Flowchart for a Hierarchical Queue



Hierarchical Queues

The switch uses a hierarchical queueing model for traffic sent from an ES port. Each packet is assigned a queue based on its physical interface, VLAN, or class:

- At the class level, a packet is queued to one of four queues according to its CoS, DSCP, IP precedence, or MPLS EXP classification. Packets can be classified by any combination of these values, but if a packet matches more than one, the classification occurs in the order listed. The last queue in each set of four queues is reserved as the default queue. Packets that are not classified into one of the other three queues are assigned to the default queue. You can configure traffic in the default queue with congestion-avoidance features and scheduling congestion-management features as described in the [“Queueing and Scheduling of Hierarchical Queues” section on page 26-26](#).
- At the VLAN level, the switch supports 2045 VLAN classes divided between the two ES ports. One queue is reserved as the default queue. Packets that are not classified into one of the other VLAN queues are assigned to the default queue. You can configure traffic in the default queue with congestion-avoidance features and scheduling congestion-management features as described in the [“Queueing and Scheduling of Hierarchical Queues” section on page 26-26](#).
- At the physical level, the switch reserves one queue per port.

The switch creates the default queue and uses it to send all traffic when an egress service policy is not attached to an ES port. User traffic on a physical port without an attached service policy bypasses the QoS classification and is queued to the default queue. The minimum and maximum bandwidth for the default queue is the same as the port bandwidth.

Under congested conditions, the switch discards packets for all classes configured for the same sending queue with equal probability. To achieve the full queueing capacity, there must be an equal division of traffic among the classes for each sending queue.

Congestion-Management and Congestion-Avoidance Features

You use congestion-management features to control congestion and to control the order in which outbound packets are sent from an ES port based on the priorities assigned to those packets. You manage congestion by creating queues, by assigning packets based on the packet classification, and by scheduling the packets to be sent from the queue.

During periods with light traffic (when no congestion exists), the switch sends packets as soon as they arrive. During periods of congestion at the outbound port, packets arrive faster than the port can send them. If you use congestion-management features, the switch queues accumulating packets at a port until it is free to send them. They are then scheduled for transmission according to their assigned priorities and the queueing mechanism for the port.

You can configure either tail drop or WRED. You cannot configure both tail drop and WRED in the same class policy, but they can be used in two different class policies in the same policy map.

You can configure CBWFQ as a queue scheduling management feature, LLQ as a scheduling congestion-management feature, and traffic shaping to decrease the burstiness of traffic.

Tail Drop

With tail drop, packets are queued for the class until the maximum threshold is exceeded, and then all the packets destined for the class queue are dropped. You enable tail drop at the class level by using the **queue-limit** policy-map class configuration command. For configuration information, see the [“Configuring an Egress Hierarchical QoS Policy” section on page 26-77](#).

WRED

Cisco Systems implements a version of Random Early Detection (RED), called WRED, differently from other congestion-avoidance techniques. WRED attempts to anticipate and avoid congestion, rather than controlling congestion when it occurs. WRED takes advantage of the TCP congestion control to try to control the average queue size by signaling end hosts when they should temporarily stop sending packets. By randomly dropping packets before periods of high congestion, it tells the packet source to decrease its sending rate. Assuming the packet source is using TCP, WRED tells it to decrease its sending rate until all the packets reach their destination, meaning that the congestion is cleared. By dropping some packets early rather than waiting until the queue is full, WRED avoids dropping large numbers of packets at once.

When a packet arrives and WRED is enabled, these events occur:

- The average queue size is calculated based on the previous average and the current size of the queue. The average queue-size calculation is affected by the exponential weight constant setting in the **random-detect exponential-weight-constant** policy-map class configuration command.
- If the average queue size is less than the minimum queue threshold, the arriving packet is queued. The minimum queue threshold is configured through the *min-threshold* option in the **random-detect {dscp | precedence}** policy-map class configuration command.
- If the average queue size is between the minimum queue threshold and the maximum queue threshold, the packet is either dropped or queued, depending on the packet-drop probability. The packet-drop probability is based on the minimum threshold, the maximum threshold, and the mark-probability denominator. The maximum queue threshold is configured through the *max-threshold* option, and the mark-probability denominator is configured through the *mark-prob-denominator* option in the **random-detect dscp {dscp | precedence}** policy-map class configuration command.
- If the average queue size is greater than the maximum queue threshold, the packet is automatically dropped.

You enable WRED by using the **random-detect** policy-map class configuration command at the class level. This command allows for preferential drop treatment among packets with different IP precedence or DSCP values. The WRED algorithm discards or marks packets destined for a queue when that queue is congested. It discards packets fairly and before the queue is full. Packets with high IP-precedence values are preferred over packets with low IP-precedence values. For configuration information, see the [“Configuring an Egress Hierarchical QoS Policy” section on page 26-77](#).

CBWFQ

CBWFQ provides guaranteed bandwidth to particular traffic classes, such as voice, that are delay sensitive, while still fairly serving all other traffic in the network. You define traffic classes based on match criteria. Packets satisfying the match criteria for a class constitute the traffic for that class. A queue is reserved for each class, and traffic belonging to a class is directed to the queue for that class.

The bandwidth assigned to a class is the minimum bandwidth that is delivered to the class during congestion. CBWFQ uses the bandwidth weight to ensure that the queue for the class is serviced fairly.

You enable CBWFQ and specify the minimum bandwidth as a rate in kbps or as a percentage of the available bandwidth by using the **bandwidth** policy-map class configuration command at the class level or at the VLAN level. During periods of congestion, the classes are serviced in proportion to their configured bandwidth. The amount of bandwidth available to a class is dependent on the amount of bandwidth reserved by the parent class. For configuration information, see the [“Configuring an Egress Hierarchical QoS Policy” section on page 26-77](#).

LLQ

LLQ provides strict-priority queueing for a traffic class. It enables delay-sensitive data, such as voice, to be sent before packets in other queues are sent. The priority queue is serviced first until it is empty. Only one traffic stream can be destined for the priority queue per class-level policy. The priority queue restricts all traffic streams in the same hierarchy, and you should use care when configuring this feature. You enable the priority queue for a traffic class by using the **priority** policy-map class configuration command at the class level. For configuration information, see the [“Configuring an Egress Hierarchical QoS Policy” section on page 26-77](#).

Shaping

Shaping provides a process for delaying out-of-profile packets in queues so that they conform to a specified profile. Shaping is distinct from policing. Policing drops packets that exceed a configured threshold, but shaping buffers packets so that traffic remains within a threshold. Shaping offers greater smoothness in handling traffic than policing. You enable average-rate traffic shaping on a traffic class by using the **shape** policy-map class configuration command at the class level or at the VLAN level. At the physical level of the hierarchy, you can shape only the class-default class by using the **shape** policy-map class configuration command in an egress policy attached to an ES port. For configuration information, see the [“Configuring an Egress Hierarchical QoS Policy” section on page 26-77](#).

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 ingress and egress queues instead of using the default QoS behavior. (The default is that QoS is disabled. The switch then 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, page 26-30](#)
- [Effects of Auto-QoS on the Configuration, page 26-33](#)
- [Auto-QoS Configuration Guidelines, page 26-33](#)
- [Enabling Auto-QoS for VoIP, page 26-34](#)
- [Auto-QoS Configuration Example, page 26-35](#)

Generated Auto-QoS Configuration

By default, auto-QoS is disabled on all ports.

When auto-QoS is enabled, it uses the ingress packet label to categorize traffic, to assign packet labels, and to configure the ingress queues and egress queue-sets as shown in [Table 26-2](#).

Table 26-2 Traffic Types, Ingress Packet Labels, Assigned Packet Labels, and Queues

	VoIP Data Traffic	VoIP Control Traffic	Routing Protocol Traffic	STP BPDU Traffic	All Other Traffic
Ingress DSCP	46	26	—	—	—
Ingress CoS	5	3	6	7	—
DiffServ	EF	AF31	—	—	—
Assigned DSCP	46	26	48	56	0
Assigned CoS	5	3	6	7	0
CoS-to-ingress queue map	2, 3, 4, 5, 6, 7 (queue 2)				0, 1 (queue 1)
CoS-to-egress queue-set map	5 (queue 1)	3, 6, 7 (queue 2)			2, 4 (queue 3) 0, 1 (queue 4)

[Table 26-3](#) shows the generated auto-QoS configuration for the ingress queues.

Table 26-3 Auto-QoS Configuration for the Ingress Queues

Ingress Queue	Queue Number	CoS-to-Queue Map	Queue Weight (Bandwidth)	Queue (Buffer) Size
SRR shared	1	0, 1	90 percent	90 percent
Priority	2	2, 3, 4, 5, 6, 7	10 percent	10 percent

[Table 26-4](#) shows the generated auto-QoS configuration for the egress queue-set.

Table 26-4 Auto-QoS Configuration for the Egress Queue-Set

Egress Queue	Queue Number in the Queue-Set	CoS-to-Queue Map	Queue Weight (Bandwidth)	Queue (Buffer) Size
Priority (shaped)	1	5	10 percent	20 percent
SRR shared	2	3, 6, 7	10 percent	20 percent
SRR shared	3	2, 4	60 percent	20 percent
SRR shared	4	0, 1	20 percent	40 percent

When you enable the auto-QoS feature on the first port, these automatic actions occur:

- QoS is globally enabled (**mls qos** global configuration command), and other global configuration commands are added.
- When you enter the **auto qos voip cisco-phone** interface configuration command on a port at the edge of the network that is connected to a Cisco IP phone, the switch enables the trusted boundary feature. The switch 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 port 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 switch configures ingress queues and the egress queue-set on the port according to the settings in [Table 26-3](#) and [Table 26-4](#).
- When you enter the **auto qos voip trust** interface configuration command on a port connected to the interior of the network, the switch trusts the CoS value for nonrouted ports or the DSCP value for routed ports in ingress packets (the assumption is that traffic has already been classified by other edge devices). The switch configures the ingress queues and the egress queue-set on the port according to the settings in [Table 26-3](#) and [Table 26-4](#).

For information about the trusted boundary feature, see the “[Configuring a Trusted Boundary to Ensure Port Security](#)” section on page 26-46.

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 the ingress packet label and applies the commands listed in [Table 26-5](#) to the port.



Note

On an ES port, the **srr-queue bandwidth shape** interface configuration command is not part of the generated **auto qos voip** command list.

Table 26-5 Generated Auto-QoS Configuration

Description	Automatically Generated Command
The switch automatically enables standard QoS and configures the CoS-to-DSCP map (maps CoS values in inbound packets to a DSCP value).	Switch(config)# mls qos Switch(config)# mls qos map cos-dscp 0 8 16 26 32 46 48 56
The switch automatically maps CoS values to an ingress queue and to a threshold ID.	Switch(config)# no mls qos srr-queue input cos-map Switch(config)# mls qos srr-queue input cos-map queue 1 threshold 3 0 Switch(config)# mls qos srr-queue input cos-map queue 1 threshold 2 1 Switch(config)# mls qos srr-queue input cos-map queue 2 threshold 1 2 Switch(config)# mls qos srr-queue input cos-map queue 2 threshold 2 4 6 7 Switch(config)# mls qos srr-queue input cos-map queue 2 threshold 3 3 5
The switch automatically maps CoS values to an egress queue in the queue-set and to a threshold ID.	Switch(config)# no mls qos srr-queue output cos-map Switch(config)# mls qos srr-queue output cos-map queue 1 threshold 3 5 Switch(config)# mls qos srr-queue output cos-map queue 2 threshold 3 3 6 7 Switch(config)# mls qos srr-queue output cos-map queue 3 threshold 3 2 4 Switch(config)# mls qos srr-queue output cos-map queue 4 threshold 2 1 Switch(config)# mls qos srr-queue output cos-map queue 4 threshold 3 0

Table 26-5 Generated Auto-QoS Configuration (continued)

Description	Automatically Generated Command
The switch automatically maps DSCP values to an ingress queue and to a threshold ID.	<pre>Switch(config)# no mls qos srr-queue input dscp-map Switch(config)# mls qos srr-queue input dscp-map queue 1 threshold 2 9 10 11 12 13 14 15 Switch(config)# mls qos srr-queue input dscp-map queue 1 threshold 3 0 1 2 3 4 5 6 7 Switch(config)# mls qos srr-queue input dscp-map queue 1 threshold 3 32 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 1 16 17 18 19 20 21 22 23 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 2 26 33 34 35 36 37 38 39 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 2 48 49 50 51 52 53 54 55 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 2 56 57 58 59 60 61 62 63 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 3 24 25 27 28 29 30 31 40 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 3 41 42 43 44 45 46 47</pre>
The switch automatically maps DSCP values to an egress queue in the queue-set and to a threshold ID.	<pre>Switch(config)# no mls qos srr-queue output dscp-map Switch(config)# mls qos srr-queue output dscp-map queue 1 threshold 3 40 41 42 43 44 45 46 47 Switch(config)# mls qos srr-queue output dscp-map queue 2 threshold 3 24 25 27 28 29 30 31 48 Switch(config)# mls qos srr-queue output dscp-map queue 2 threshold 3 49 50 51 52 53 54 55 56 Switch(config)# mls qos srr-queue output dscp-map queue 2 threshold 3 57 58 59 60 61 62 63 Switch(config)# mls qos srr-queue output dscp-map queue 3 threshold 3 16 17 18 19 20 21 22 23 Switch(config)# mls qos srr-queue output dscp-map queue 3 threshold 3 26 32 33 34 35 36 37 38 Switch(config)# mls qos srr-queue output dscp-map queue 3 threshold 3 39 Switch(config)# mls qos srr-queue output dscp-map queue 4 threshold 1 8 Switch(config)# mls qos srr-queue output dscp-map queue 4 threshold 2 9 10 11 12 13 14 15 Switch(config)# mls qos srr-queue output dscp-map queue 4 threshold 3 0 1 2 3 4 5 6 7</pre>
The switch automatically sets up the ingress queues, with queue 2 as the priority queue and queue 1 in shared mode. The switch also configures the bandwidth and buffer size for the ingress queues.	<pre>Switch(config)# no mls qos srr-queue input priority-queue 1 Switch(config)# no mls qos srr-queue input priority-queue 2 Switch(config)# mls qos srr-queue input bandwidth 90 10 Switch(config)# no mls qos srr-queue input buffers</pre>
The switch automatically configures the egress queue-set buffer sizes. It configures the bandwidth and the SRR mode (shaped or shared).	<pre>Switch(config)# mls qos queue-set output 1 buffers 20 20 20 40 Switch(config-if)# srr-queue bandwidth shape 10 0 0 0 Switch(config-if)# srr-queue bandwidth share 10 10 60 20</pre>

Table 26-5 Generated Auto-QoS Configuration (continued)

Description	Automatically Generated Command
The switch automatically sets the ingress classification to trust the CoS value received in the packet on a nonrouted port or to trust the DSCP value received in the packet on a routed port.	Switch(config-if)# mls qos trust cos Switch(config-if)# mls qos trust dscp
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

Effects of Auto-QoS on the Configuration

When you enable auto-QoS, the switch adds the **auto qos voip** interface configuration command and the generated configuration to the running configuration.

The switch applies the auto-QoS-generated commands as if the commands were entered from the CLI. An existing user configuration can cause the application of the generated commands to fail, or the user configuration might be overridden by the generated commands. These actions occur without warning. If all the generated commands are successfully applied, any user-entered configuration that was not overridden remains in the running configuration. Any user-entered configuration that was overridden can be retrieved by reloading the switch without saving the current configuration to memory. If the generated commands are not applied, the previous running configuration is restored.

Auto-QoS 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, you should enable auto-QoS before you configure other 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. For more information, see the [“Effects of Auto-QoS on the Configuration”](#) section on page 26-33.
- You can enable auto-QoS on static, dynamic-access, voice VLAN access, and trunk ports.
- By default, the CDP is enabled on all ports. For auto-QoS to function properly, do not disable the CDP.
- Policing is not enabled with auto-QoS. You can manually enable policing, as described in the [“Configuring an Ingress QoS Policy”](#) section on page 26-48.

Enabling Auto-QoS for VoIP

Beginning in privileged EXEC mode, follow these steps to enable auto-QoS for VoIP within a QoS domain:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port that is connected to a Cisco IP phone or the uplink port that is connected to another switch or router in the interior of the network, and enter interface configuration mode.
Step 3	auto qos voip { cisco-phone trust }	Enable auto-QoS. The keywords have these meanings: • cisco-phone—If the port is connected to a Cisco IP phone, the QoS labels of inbound packets are trusted only when the telephone is detected. • trust—The uplink port is connected to a trusted switch or router, and the VoIP traffic classification in the ingress packet is trusted.
Step 4	end	Return to privileged EXEC mode.
Step 5	show auto qos interface <i>interface-id</i>	Verify your entries. This command displays the initial auto-QoS configuration that was applied; it does not display any user changes to the configuration that might be in effect. You can use the show running-config privileged EXEC command to display the auto-QoS configuration and the user modifications.

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, refer to the “debug autoqos” command in the command reference for this release.

To disable auto-QoS on a port, use the **no auto qos voip** interface configuration command. Only the auto-QoS-generated interface configuration commands for this port are removed. If this is the last port on which auto-QoS is enabled and you enter the **no auto qos voip** command, auto-QoS is considered disabled even though the auto-QoS-generated global configuration commands remain (to avoid disrupting traffic on other ports affected by the global configuration). You can use the **no mls qos** global configuration command to disable the auto-QoS-generated global configuration commands. With QoS disabled, there is no concept of trusted or untrusted ports because the packets are not modified (the CoS, DSCP, and IP precedence values in the packet are not changed). Traffic is switched in pass-through mode (packets are switched without any rewrites and classified as best effort without any policing).


This example shows how to enable auto-QoS on a port and to trust the QoS labels received in inbound packets when the switch or router connected to a port is a trusted device:

```
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# auto qos voip trust
```

Auto-QoS Configuration Example

This section describes how you could implement auto-QoS in a network, as shown in [Figure 26-12](#).

Figure 26-12 Auto-QoS Configuration Example Network



[Figure 26-12](#) shows a network in which the VoIP traffic is prioritized over all other traffic. Auto-QoS is enabled on the switches in the wiring closets at the edge of the QoS domains.



Note

You should not configure any standard QoS commands before entering the auto-QoS commands. You can fine-tune the QoS configuration, but we recommend that you do so only after the auto-QoS configuration is completed.

Beginning in privileged EXEC mode, follow these steps to configure the switch at the edge of the QoS domain to prioritize the VoIP traffic over all other traffic:

	Command	Purpose
Step 1	debug autoqos	Enable debugging for auto-QoS. When debugging is enabled, the switch displays the QoS configuration that is automatically generated when auto-QoS is enabled.
Step 2	configure terminal	Enter global configuration mode.
Step 3	cdp enable	Enable CDP globally. By default, it is enabled.
Step 4	interface interface-id	Specify the switch port connected to the Cisco IP phone, and enter interface configuration mode.
Step 5	auto qos voip cisco-phone	Enable auto-QoS on the port, and specify that the port is connected to a Cisco IP phone. The QoS labels of inbound packets are trusted only when the Cisco IP phone is detected.
Step 6	exit	Return to global configuration mode.
Step 7		Repeat Steps 4 to 6 for as many ports as are connected to the Cisco IP phone.
Step 8	auto qos voip cisco-phone	Enable auto-QoS on the port, and specify that the port is connected to a Cisco IP phone.
Step 9	exit	Return to global configuration mode.
Step 10	interface interface-id	Specify the port identified as connected to a trusted switch or router. See Figure 26-12 . Enter interface configuration mode.
Step 11	auto qos voip trust	Enable auto-QoS on the port, and specify that the port is connected to a trusted router or switch.
Step 12	end	Return to privileged EXEC mode.
Step 13	show auto qos	Verify your entries. This command displays the auto-QoS configuration that is initially applied; it does not display any user changes to the configuration that might be in effect. For information about the QoS configuration that might be affected by auto-QoS, see the “Displaying Auto-QoS Information” section on page 26-12.
Step 14	copy running-config startup-config	Save the auto qos voip interface configuration commands and the generated auto-QoS configuration in the configuration file.

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 displays 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 maps cos-dscp**
- **show mls qos interface [interface-id] [buffers | queueing]**
- **show mls qos maps [cos-dscp | cos-input-q | cos-output-q | dscp-cos | dscp-input-q | dscp-output-q]**
- **show mls qos input-queue**
- **show running-config**

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.

These sections describe how to configure standard QoS on your switch:

- [Default Standard QoS Configuration, page 26-38](#)
- [Standard QoS Configuration Guidelines, page 26-40](#)
- [Packet Modification, page 26-41](#)
- [Enabling QoS Globally, page 26-42](#) (required)
- [Configuring Ingress Classification by Using Port Trust States, page 26-42](#) (required)
- [Configuring an Ingress QoS Policy, page 26-48](#) (required)
- [Configuring DSCP Maps, page 26-59](#) (optional, unless you need to use the DSCP-to-DSCP-mutation map or the policed-DSCP map)
- [Configuring Ingress Queue Characteristics, page 26-64](#) (optional)
- [Configuring Egress Queue-Set Characteristics, page 26-69](#) (optional)

If you need to configure outbound traffic on an ES port, see the [“Configuring Hierarchical QoS” section on page 26-76](#).

Default Standard QoS Configuration

QoS is disabled. There is no concept of trusted or untrusted ports because the packets are not modified (the CoS, DSCP, and IP precedence values in the packet are not changed). Traffic is switched in pass-through mode (packets are switched without any rewrites and classified as best effort without any policing).

When QoS is enabled with the **mls qos** global configuration command and all other QoS settings are at their defaults, traffic is classified as best effort (the DSCP and CoS value is set to 0) without any policing. No policy maps are configured. The default port trust state on all ports is untrusted. The default ingress queues and egress queue-set settings are described in the [“Default Ingress Queue Configuration” section on page 26-38](#) and the [“Default Egress Queue-Set Configuration” section on page 26-39](#). For outbound traffic on an ES port, see the [“Default Hierarchical QoS Configuration” section on page 26-76](#).

Default Ingress Queue Configuration

[Table 26-6](#) shows the default ingress queue configuration when QoS is enabled.

Table 26-6 Default Ingress Queue Configuration

Feature	Queue 1	Queue 2
Buffer allocation	90 percent	10 percent
Bandwidth allocation ¹	4	4
Priority queue bandwidth ²	0	10
WTD drop threshold 1	100 percent	100 percent
WTD drop threshold 2	100 percent	100 percent

1. The bandwidth is equally shared between the queues. SRR sends packets in shared mode only.

2. Queue 2 is the priority queue. SRR services the priority queue for its configured share before servicing the other queue.

[Table 26-7](#) shows the default CoS input queue threshold map when QoS is enabled.

Table 26-7 Default CoS Input Queue Threshold Map

CoS Value	0–4	5	6, 7
Queue ID - Threshold ID	1 - 1	2 - 1	1 - 1

[Table 26-8](#) shows the default DSCP input queue threshold map when QoS is enabled.

Table 26-8 Default DSCP Input Queue Threshold Map

DSCP Value	0–39	40–47	48–63
Queue ID - Threshold ID	1 - 1	2 - 1	1 - 1

Default Egress Queue-Set Configuration

Table 26-9 shows the default egress queue-set configuration when QoS is enabled. All ports are mapped to queue-set 1. The port bandwidth limit is set to 100 percent, and the rate is unlimited.

Table 26-9 Default Egress Queue-Set Configuration

Feature	Queue 1	Queue 2	Queue 3	Queue 4
Buffer allocation	25 percent	25 percent	25 percent	25 percent
WTD drop threshold 1	100 percent	50 percent	100 percent	100 percent
WTD drop threshold 2	100 percent	50 percent	100 percent	100 percent
Reserved threshold	50 percent	100 percent	50 percent	50 percent
Maximum threshold	400 percent	400 percent	400 percent	400 percent
SRR shaped weights (absolute) ¹	25	0	0	0
SRR shared weights ²	25	25	25	25

1. A shaped weight of zero means that this queue is operating in shared mode.
2. One quarter of the bandwidth is allocated to each queue.

Table 26-10 shows the default CoS output queue threshold map when QoS is enabled.

Table 26-10 Default CoS Output Queue Threshold Map

CoS value	0, 1	2, 3	4	5	6, 7
Queue ID - threshold ID	2 - 1	3 - 1	4 - 1	1 - 1	4 - 1

Table 26-11 shows the default DSCP output queue threshold map when QoS is enabled.

Table 26-11 Default DSCP Output Queue Threshold Map

DSCP Value	0–15	16–31	32–39	40–47	48–63
Queue ID - Threshold ID	2 - 1	3 - 1	4 - 1	1 - 1	4 - 1

Default Mapping Table Configuration

The default CoS-to-DSCP map is shown in Table 26-12 on page 26-59.

The default IP-precedence-to-DSCP map is shown in Table 26-13 on page 26-60.

The default DSCP-to-CoS map is shown in Table 26-14 on page 26-62.

The default DSCP-to-DSCP-mutation map is a null map, which maps an inbound DSCP value to the same DSCP value.

The default policed-DSCP map is a null map, which maps an inbound DSCP value to the same DSCP value (no markdown).

Standard QoS Configuration Guidelines

Before beginning the QoS configuration, you should be aware of this information:

- It is not possible to match IP fragments against configured IP extended ACLs to enforce QoS. IP fragments are sent as best-effort. IP fragments are denoted by fields in the IP header.
- Only one ACL can be configured per class map. The ACL can have multiple ACEs, which match fields against the contents of the packet. Class maps that contain ACLs are supported only in an ingress policy-map.
- Only one policy map per port is supported. You can attach one ingress service-policy per port and one egress service-policy per ES port.
- Inbound traffic is classified, policed, and marked down (if configured) regardless of whether the traffic is bridged, routed, or sent to the CPU. Bridged frames can be dropped or have their DSCP and CoS values modified.
- Only one ingress policer is applied to a packet on a port. Only the average-rate and committed-burst parameters are configurable.
- You can create an aggregate policer that is shared by multiple traffic classes within the same policy map. However, you cannot use the aggregate policer across different policy maps.
- On a standard port and on the input of an ES port, the port ASIC device, which controls more than one physical port, supports 256 ingress policers (255 policers plus 1 **no** policer). The maximum number of policers supported per is 64. For example, you could configure 32 policers on a Gigabit Ethernet port and 8 policers on a Fast Ethernet port, or you could configure 64 policers on a Gigabit Ethernet port and 5 policers on a Fast Ethernet port. Policers are allocated on demand by the software and are constrained by the hardware and ASIC boundaries. You cannot reserve policers per port; there is no guarantee that a port will be assigned to any policer. These limitations do not apply to egress policers configured on the ES ports.
- On a port configured for QoS, all traffic received through the port is classified, policed, and marked according to the policy map attached to the port. On a trunk port configured for QoS, traffic in *all* VLANs received through the port is classified, policed, and marked according to the policy map attached to the port.
- Because the switch does not support attaching a service policy to a logical interface (such as an EtherChannel), 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.
- Control traffic (such as spanning-tree bridge protocol data units [BPDUs] and routing update packets) received by the switch are subject to all ingress QoS processing.
- You are likely to lose data when you change queue settings; therefore, try to make changes when traffic is at a minimum.

For outbound traffic on an ES port, see the [“Hierarchical QoS Configuration Guidelines” section on page 26-76](#).

Packet Modification

A packet is classified, policed, and queued to provide QoS. Packet modifications can occur during this process:

- For IP and non-IP packets, ingress classification involves assigning a QoS label to a packet based on the DSCP or CoS of the received packet. However, the packet is not modified at this stage; only an indication of the assigned DSCP or CoS value is carried along. The reason for this is that QoS classification and forwarding lookups occur in parallel, and it is possible that the packet is forwarded with its original DSCP to the CPU where it is again processed through software.

The ES ports classify traffic on egress, and this classification can be used for queuing or for marking the CoS, DSCP, IP precedence, or MPLS EXP bits. Any packet modifications that result from ingress classification are applied before the packet reaches the egress classification stage. For example, if the switch receives traffic with a CoS value of 2 and an ingress action resets the CoS to 4, the packet will have a CoS of 4 (instead of a CoS of 2 and an indicator that the CoS should be set to 4) when it moves to the egress classification stage.

- During ingress policing, IP and non-IP packets can have another DSCP assigned to them (if they are out of profile and the policer specifies a markdown DSCP). Once again, the DSCP in the packet is not modified, but an indication of the marked-down value is carried along. For IP packets, the packet modification occurs at a later stage; for non-IP packets the DSCP is converted to CoS and used for queueing and scheduling decisions.

During egress policing on the ES ports, marking actions can set the CoS, DSCP, IP precedence, or the MPLS EXP bits. Any markings performed by an ingress policer are applied before the packet reaches the egress classification stage.

- Depending on the QoS label assigned to a frame and the mutation chosen, the DSCP and CoS values of the frame are rewritten. If you do not configure the mutation map and if you configure the port to trust the DSCP of the inbound frame, the DSCP value in the frame is not changed, but the CoS is rewritten according to the DSCP-to-CoS map. If you configure the port to trust the CoS of the inbound frame and it is an IP packet, the CoS value in the frame is not changed, but the DSCP might be changed according to the CoS-to-DSCP map.

The input mutation causes the DSCP to be rewritten depending on the new value of DSCP chosen. The set action in a policy map also causes the DSCP to be rewritten.

This information applies to both standard and ES ports. On the ES ports, the switch also applies trust policies to 802.1Q tunneling frames at egress.

Enabling QoS Globally

By default, QoS is disabled on the switch.

Beginning in privileged EXEC mode, follow these steps to enable QoS. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos	Enable QoS globally. QoS runs from the default settings described in the “Default Standard QoS Configuration” section on page 26-38 and the “Default Hierarchical QoS Configuration” section on page 26-76.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable QoS, use the **no mls qos** global configuration command.

Configuring Ingress Classification by Using Port Trust States


These sections describe how to classify inbound traffic through port trust states. Depending on your network configuration, you must perform one or more of these tasks or one or more of the tasks in the [“Configuring an Ingress QoS Policy”](#) section on page 26-48:

- [Configuring the Trust State on Ports Within the QoS Domain, page 26-42](#)
- [Configuring the CoS Value for an Interface, page 26-45](#)
- [Configuring a Trusted Boundary to Ensure Port Security, page 26-46](#)
- [Configuring the DSCP Trust State on a Port Bordering Another QoS Domain, page 26-47](#)

Configuring the Trust State on Ports Within the QoS Domain

Packets entering a QoS domain are classified at the edge of the QoS domain. The switch port within the QoS domain can then be configured to one of the trusted states because there is no need to classify the packets at every switch within the domain. [Figure 26-13](#) shows a sample network topology.

Figure 26-13 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:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface interface-id	Specify the port to be trusted, and enter interface configuration mode. Valid interfaces include physical ports.

	Command	Purpose
Step 3	mls qos trust [cos dscp ip-precedence]	Configure the port trust state. By default, the port is not trusted. If no keyword is specified, the default is dscp. For 802.1Q tunnels, the switch processes inbound traffic on a standard port according to the trusted setting applied to this port. The switch configures the inner and outer tags for packets sent over the ES trunk port. The keywords have these meanings: • cos—Classifies an ingress packet by using the packet CoS value. For an untagged packet, the port default CoS value is used. The default port CoS value is 0. For 802.1Q tunnels, the switch copies the inner CoS value to the outer CoS value and sends the packet out an ES port. • dscp—Classifies an ingress packet by using the packet DSCP value if the packet is an IP packet. For a non-IP packet, the packet CoS value is used if the packet is tagged; for an untagged packet, the default port CoS is used. Internally, the switch maps the CoS value to a DSCP value by using the CoS-to-DSCP map. For 802.1Q tunnels, for a non-IP packet that is untagged, the switch configures the outer CoS value from the DSCP-to-CoS map, does not modify the inner CoS value, and sends the packet out an ES port. For an IP packet, the switch modifies the DSCP value in the packet if there is a DSCP-to-DSCP mutation map configured on the standard port. The switch uses the mutated DSCP value to configure the outer CoS value from the DSCP-to-CoS map and sends the packet out an ES port. • ip-precedence—Classifies an ingress packet by using the packet IP-precedence value. For a non-IP packet, the packet CoS value is used if the packet is tagged; for an untagged packet, the default port CoS is used. Internally, the switch maps the CoS value to a DSCP value through the CoS-to-DSCP map. For 802.1Q tunnels, the switch converts the generated DSCP value from the DSCP-to-CoS map and uses it as the outer CoS value in the packet. The switch does not modify the inner CoS value in the packet and sends the packet out an ES port. Note When port trust policies are used with 802.1Q tunneling, all ports sharing the same tunnel VLAN must be configured with the same trust policy, and the ports involved must use the same DSCP-to-DSCP mutation map. For more information, see the “Configuring the DSCP-to-DSCP-Mutation Map” section on page 26-63.
Step 4	end	Return to privileged EXEC mode.
Step 5	show mls qos interface	Verify your entries.
Step 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 26-45](#). For information on how to configure the CoS-to-DSCP map, see the [“Configuring the CoS-to-DSCP Map” section on page 26-59](#).

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 inbound packets on the port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port to be configured, and enter interface configuration mode. Valid interfaces include physical ports.
Step 3	mls qos cos { <i>default-cos</i> override }	Configure the default CoS value for the port. For 802.1Q tunnels, the switch processes inbound traffic on a standard port according to the trusted setting applied to this port. The switch configures the inner and outer tags for packets sent over the ES trunk port. - For <i>default-cos</i>, specify a default CoS value to be assigned to a port. If the packet is untagged, the default CoS value becomes the packet CoS value. The CoS range is 0 to 7. The default is 0. For 802.1Q tunnels, the switch copies the inner CoS value to the outer CoS value and sends the packet out an ES port. - Use the override keyword to override the previously configured trust state of the inbound packet and to apply the default port CoS value to the port on all inbound packets. By default, CoS override is disabled. Use the override keyword when all inbound packets on specified ports deserve higher or lower priority than packets entering from other ports. Even if a port was previously set to trust DSCP, CoS, or IP precedence, this command overrides the previously configured trust state, and all the inbound CoS values are assigned the default CoS value configured with this command. If an inbound packet is tagged, the CoS value of the packet is modified with the default CoS of the port at the ingress port. For 802.1Q tunnels, the switch copies the port CoS value to the outer CoS value, does not modify the inner CoS value, and sends the packet out an ES port.
Step 4	end	Return to privileged EXEC mode.
Step 5	show mls qos interface	Verify your entries.
Step 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 a Trusted Boundary to Ensure Port Security

In a typical network, you connect a Cisco IP phone to a switch port, as shown in [Figure 26-13 on page 26-43](#), and cascade devices that generate data packets from the back of the telephone. The Cisco IP phone guarantees the voice quality through a shared data link by marking the CoS level of the voice packets as high priority (CoS = 5) and by marking the data packets as low priority (CoS = 0). 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 is the priority of the packet.

For most Cisco IP phone configurations, the traffic sent from the telephone to the switch should be 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 configure the switch port to which the telephone is connected to trust the CoS labels of all traffic received on that port.

With the trusted setting, you also can use the trusted boundary feature to prevent misuse of a high-priority queue if a user bypasses the telephone and connects the PC directly to the switch. Without trusted boundary, the CoS labels generated by the PC are trusted by the switch (because of the trusted CoS setting). By contrast, trusted boundary uses 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. Note that the trusted boundary feature is not effective if the PC and Cisco IP phone are connected to a hub that is connected to the switch.

In some situations, you can prevent a PC connected to the IP phone from taking advantage of a high-priority data queue. 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.

Beginning in privileged EXEC mode, follow these steps to enable trusted boundary on a port:


	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	cdp run	Enable CDP globally. By default, CDP is enabled.
Step 3	interface <i>interface-id</i>	Specify the port connected to the IP phone, and enter interface configuration mode. Valid interfaces include physical ports.
Step 4	cdp enable	Enable CDP on the port. By default, CDP is enabled.
Step 5	mls qos trust cos	Configure the port to trust the CoS value in traffic received from the Cisco IP phone. By default, the port is not trusted.
Step 6	mls qos trust device cisco-phone	Specify that the Cisco IP phone is a trusted device. You cannot enable both trusted boundary and auto-QoS (auto qos voip interface configuration command) at the same time; they are mutually exclusive.
Step 7	end	Return to privileged EXEC mode.
Step 8	show mls qos interface	Verify your entries.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the trusted boundary feature, use the **no mls qos trust device** interface configuration command.

Configuring the DSCP Trust State on a Port Bordering Another QoS Domain

If you are administering two separate QoS domains between which you want to implement QoS features for IP traffic, you can configure the switch ports bordering the domains to a DSCP-trusted state as shown in Figure 26-14. Then the receiving port accepts the DSCP-trusted value and avoids the classification stage of QoS. If the two domains use different DSCP values, you can configure the DSCP-to-DSCP-mutation map to translate a set of DSCP values to match the definition in the other domain.

Figure 26-14 DSCP-Trusted State on a Port Bordering Another QoS Domain



Beginning in privileged EXEC mode, follow these steps to configure the DSCP-trusted state on a port and modify the DSCP-to-DSCP-mutation map. To ensure a consistent mapping strategy across both QoS domains, you must perform this procedure on the ports in both domains:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos map dscp-mutation <i>dscp-mutation-name in-dscp to out-dscp</i>	Modify the DSCP-to-DSCP-mutation map. The default DSCP-to-DSCP-mutation map is a null map, which maps an inbound DSCP value to the same DSCP value. - For <i>dscp-mutation-name</i>, enter the mutation map name. You can create more than one map by specifying a new name. - For <i>in-dscp</i>, enter up to eight DSCP values separated by spaces. Then enter the to keyword. - For <i>out-dscp</i>, enter a single DSCP value. The DSCP range is 0 to 63.
Step 3	interface interface-id	Specify the port to be trusted, and enter interface configuration mode. Valid interfaces include physical ports.
Step 4	mls qos trust dscp	Configure the ingress port as a DSCP-trusted port. By default, the port is not trusted.

	Command	Purpose
Step 5	mls qos dscp-mutation <i>dscp-mutation-name</i>	Apply the map to the specified ingress DSCP-trusted port. For <i>dscp-mutation-name</i> , specify the mutation map name created in Step 2. You can configure multiple DSCP-to-DSCP-mutation maps on an ingress port.
Step 6	end	Return to privileged EXEC mode.
Step 7	show mls qos maps dscp-mutation	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return a port to its non-trusted state, use the **no mls qos trust** interface configuration command. To return to the default DSCP-to-DSCP-mutation map values, use the **no mls qos map dscp-mutation dscp-mutation-name** global configuration command.

This example shows how to configure a port to the DSCP-trusted state and to modify the DSCP-to-DSCP-mutation map (named *gi1/1/1-mutation*) so that inbound DSCP values 10 to 13 are mapped to DSCP 30:

```
Switch(config)# mls qos map dscp-mutation gi1/1/1-mutation 10 11 12 13 to 30
Switch(config)# interface gigabitethernet1/1/1
Switch(config-if)# mls qos trust dscp
Switch(config-if)# mls qos dscp-mutation gi1/1/1-mutation
Switch(config-if)# end
```

Configuring an Ingress QoS Policy

Configuring an ingress QoS policy typically requires classifying traffic into classes, configuring policies applied to those traffic classes, and attaching the policies to a port.

For background information, see the [“Ingress Classification” section on page 26-6](#) and the [“Ingress Policing and Marking” section on page 26-9](#). For configuration guidelines, see the [“Standard QoS Configuration Guidelines” section on page 26-40](#).

These sections describe how to classify, police, and mark inbound traffic. Depending on your network configuration, you must perform one or more of these tasks:

- [Classifying Ingress Traffic by Using ACLs, page 26-49](#)
- [Classifying Ingress Traffic by Using Class Maps, page 26-52](#)
- [Classifying, Policing, and Marking Ingress Traffic by Using Policy Maps, page 26-54](#)
- [Classifying, Policing, and Marking Ingress Traffic by Using Aggregate Policers, page 26-57](#)

For information on configuring egress policies for the ES ports, see the [“Configuring an Egress Hierarchical QoS Policy” section on page 26-77](#). This section describes how to classify egress traffic by using class maps, how to configure an egress two-rate traffic policer, how to configure class-based packet marking in an egress traffic policy, how to configure CBWFQ, tail drop, DSCP-based WRED, and IP precedence-based WRED, how to enable LLQ, and how to configure shaping.

Classifying Ingress Traffic by Using ACLs

You can classify ingress IP traffic by using IP standard or IP extended ACLs. You also can classify ingress non-IP traffic by using Layer 2 MAC ACLs.

Beginning in privileged EXEC mode, follow these steps to create an IP standard ACL for inbound IP traffic:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	access-list <i>access-list-number</i> { deny permit } <i>source</i> [<i>source-wildcard</i>]	Create an IP standard ACL, repeating the command as many times as necessary. - For <i>access-list-number</i>, enter the access list number. The range is 1 to 99 and 1300 to 1999. - Use the permit keyword to permit a certain type of traffic if the conditions are matched. Use the deny keyword to deny a certain type of traffic if conditions are matched. - For <i>source</i>, enter the network or host from which the packet is being sent. You can use the any keyword as an abbreviation for 0.0.0.0 255.255.255.255. (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. 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show access-lists	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete an access list, use the **no access-list** *access-list-number* global configuration command.

This example shows how to allow access for only those hosts on the three 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 access list statements is rejected.

```
Switch(config)# access-list 1 permit 192.5.255.0 0.0.0.255
Switch(config)# access-list 1 permit 128.88.0.0 0.0.255.255
Switch(config)# access-list 1 permit 36.0.0.0 0.0.0.255
! (Note: all other access implicitly denied)
```

Beginning in privileged EXEC mode, follow these steps to create an IP extended ACL for inbound IP traffic:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	access-list <i>access-list-number</i> { deny permit } <i>protocol source source-wildcard destination destination-wildcard</i>	Create an IP extended ACL, repeating the command as many times as necessary. - For <i>access-list-number</i>, enter the access list number. The range is 100 to 199 and 2000 to 2699. - Use the permit keyword to permit a certain type of traffic if the conditions are matched. Use the deny keyword to deny a certain type of traffic if conditions are matched. - 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. - For <i>source</i>, enter the network or host from which the packet is being sent. You specify this 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. - For <i>source-wildcard</i>, enter the wildcard bits by placing ones in the bit positions that you want to ignore. You specify the wildcard 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. - 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 by <i>source</i> and <i>source-wildcard</i>. 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.
Step 3	end	Return to privileged EXEC mode.
Step 4	show access-lists	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete an access list, use the **no access-list** *access-list-number* global configuration command.

This example shows how to create an ACL that permits IP traffic from any source to any destination that has the DSCP value set to 32:

```
Switch(config)# access-list 100 permit ip any any dscp 32
```

This example shows how to create an ACL that permits IP traffic from a source host at 10.1.1.1 to a destination host at 10.1.1.2 with a precedence value of 5:

```
Switch(config)# access-list 100 permit ip host 10.1.1.1 host 10.1.1.2 precedence 5
```

This example shows how to create an ACL that permits PIM traffic from any source to a destination group address of 224.0.0.2 with a DSCP set to 32:

```
Switch(config)# access-list 102 permit pim any 224.0.0.2 dscp 32
```

Beginning in privileged EXEC mode, follow these steps to create a Layer 2 MAC ACL for inbound non-IP traffic:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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.
Step 3	{permit deny} {host <i>src-MAC-addr mask</i> any host <i>dst-MAC-addr</i> <i>dst-MAC-addr mask</i> } [<i>type mask</i>]	Specify the type of traffic to permit or deny if the conditions are matched, entering the command as many times as necessary. - For <i>src-MAC-addr</i>, enter the MAC address of the host from which the packet is being sent. You specify this by using the hexadecimal format (H.H.H), by using the any keyword as an abbreviation for <i>source</i> 0.0.0, <i>source-wildcard</i> 255.255.255, or by using the host keyword for <i>source</i> 0.0.0. - For <i>mask</i>, enter the wildcard bits by placing ones in the bit positions that you want to ignore. - For <i>dst-MAC-addr</i>, enter the MAC address of the host to which the packet is being sent. You specify this by using the hexadecimal format (H.H.H), by using the any keyword as an abbreviation for <i>source</i> 0.0.0, <i>source-wildcard</i> 255.255.255, or by using the host keyword for <i>source</i> 0.0.0. (Optional) For <i>type mask</i>, specify the Ethertype number of a packet with Ethernet II or SNAP encapsulation to identify the protocol of the packet. For <i>type</i>, the range is from 0 to 65535, typically specified in hexadecimal. For <i>mask</i>, enter the <i>don't care</i> bits applied to the Ethertype before testing for a match. 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show access-lists [<i>access-list-number</i> <i>access-list-name</i>]	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete an access list, use the **no mac access-list extended** *access-list-name* global configuration command.

This example shows how to create a Layer 2 MAC ACL with two permit statements. The first statement allows traffic from the host with MAC address 0001.0000.0001 to the host with MAC address 0002.0000.0001. The second statement allows only EtherType XNS-IDP traffic from the host with MAC address 0001.0000.0002 to the host with MAC address 0002.0000.0002.

```
Switch(config)# mac access-list extended maclist1
Switch(config-ext-macl)# permit 0001.0000.0001 0.0.0 0002.0000.0001 0.0.0
Switch(config-ext-macl)# permit 0001.0000.0002 0.0.0 0002.0000.0002 0.0.0 xns-idp
! (Note: all other access implicitly denied)
```

Classifying Ingress Traffic by Using Class Maps

You use the **class-map** global configuration command to create a class map for matching packets to the class whose name you specify. The class map isolates a specific ingress traffic flow (class) from all other traffic by defining the criteria to use to match against a specific flow. A match criterion is defined with a match statement entered within the class-map configuration mode. Packets are checked against the match criteria configured for a class map. If a packet matches the specified criteria, the packet is considered a member of the class and is forwarded according to the QoS specifications set in the traffic policy.



Note

You cannot use ACL matches in service policies attached to ES ports.

For information on how to classify egress traffic on an ES port, see the [“Classifying Egress Traffic by Using Class Maps” section on page 26-78](#).

Beginning in privileged EXEC mode, follow these steps to create a class map and to define the match criterion to classify inbound traffic:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	access-list access-list-number {deny permit} source [source-wildcard] or access-list access-list-number {deny permit} protocol source [source-wildcard] destination [destination-wildcard] or mac access-list extended name {permit deny} {host src-MAC-addr mask any host dst-MAC-addr dst-MAC-addr mask} [type mask]	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 “Classifying Ingress Traffic by Using ACLs” section on page 26-49 . 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.

	Command	Purpose
Step 3	class-map [match-all match-any] <i>class-map-name</i>	Create a class map, and enter class-map configuration mode. By default, no class maps are defined. (Optional) Use the match-all keyword to perform a logical-AND of all matching statements under this class map. All criteria in the class map must be matched. (Optional) Use the match-any keyword to perform a logical-OR of all matching statements under this class map. One or more criteria must be matched. - For <i>class-map-name</i>, specify the name of the class map. If neither the match-all nor the match-any keyword is specified, the default is match-all.
Step 4	match { access-group <i>acl-index-or-name</i> ip dscp <i>dscp-list</i> ip precedence <i>ip-precedence-list</i> }	Define the match criterion to classify traffic. By default, no match criterion is defined. Only one ACL per class map is supported. - For access-group <i>acl-index-or-name</i>, specify the number or name of the ACL created in Step 2. - For ip dscp <i>dscp-list</i>, enter a list of up to eight IP DSCP values to match against inbound packets. Separate each value with a space. The range is 0 to 63. - For ip precedence <i>ip-precedence-list</i>, enter a list of up to eight IP-precedence values to match against inbound packets. Separate each value with a space. The range is 0 to 7.
Step 5	end	Return to privileged EXEC mode.
Step 6	show class-map	Verify your entries.
Step 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** [**match-all** | **match-any**] *class-map-name* global configuration command. To remove a match criterion, use the **no match** {**access-group** *acl-index-or-name* | **ip dscp** | **ip precedence**} class-map configuration command.

This example shows how to configure the class map called *class1*. The *class1* has one match criterion, which is access list 103. It permits traffic from any host to any destination that matches a DSCP value of 10.

```
Switch(config)# access-list 103 permit any any dscp 10
Switch(config)# class-map class1
Switch(config-cmap)# match access-group 103
Switch(config-cmap)# exit
```

This example shows how to create a class map called *class2*, which matches inbound traffic with DSCP values of 10, 11, and 12:

```
Switch(config)# class-map class2
Switch(config-cmap)# match ip dscp 10 11 12
Switch(config-cmap)# exit
```

This example shows how to create a class map called *class3*, which matches inbound traffic with IP-precedence values of 5, 6, and 7:

```
Switch(config)# class-map class3
Switch(config-cmap)# match ip precedence 5 6 7
Switch(config-cmap)# exit
```

Classifying, Policing, and Marking Ingress Traffic by Using Policy Maps

A policy map specifies which inbound traffic class to act on. You can specify which CoS, DSCP, or IP precedence values in the traffic class to trust. You can specify which DSCP or IP precedence values in the traffic class to set. You can specify the traffic bandwidth limitations for each matched traffic class (policer) and the action to take when the traffic is out of profile (marking).

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 a port.
- A policy-map trust state and a port trust state are mutually exclusive. The last one configured takes affect.

You can attach only one ingress policy-map per port.

Before beginning this procedure, make sure that you have created the class map to isolate traffic. For more information, see the [“Classifying Ingress Traffic by Using ACLs” section on page 26-49](#) and the [“Classifying Ingress Traffic by Using Class Maps” section on page 26-52](#).

For information about configuring an egress policy-map for an ES port, see the [“Configuring an Egress Two-Rate Traffic Policer” section on page 26-80](#) and the [“Configuring Class-Based Packet Marking in an Egress Traffic Policy” section on page 26-84](#).

Beginning in privileged EXEC mode, follow these steps to create an ingress policy-map:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	policy-map <i>policy-map-name</i>	Create a policy map by entering the policy-map name, and enter policy-map configuration mode. By default, no policy maps are defined. 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.
Step 3	class <i>class-name</i>	Specify the name of the class whose traffic policy you want to create or change, and enter policy-map class configuration mode. By default, no traffic classes are defined.

	Command	Purpose
Step 4	trust [cos dscp ip-precedence]	Configure the trust state, which QoS uses to generate a CoS-based or DSCP-based QoS label. By default, the port is not trusted. If no keyword is specified when the command is entered, the default is dscp. The keywords have these meanings: • cos—QoS derives the DSCP value by using the received or default port CoS value and the CoS-to-DSCP map. • dscp—QoS derives the DSCP value by using the DSCP value from the ingress packet. For non-IP packets that are tagged, QoS derives the DSCP value by using the received CoS value; for non-IP packets that are untagged, QoS derives the DSCP value by using the default port CoS value. In either case, the DSCP value is derived from the CoS-to-DSCP map. • ip-precedence—QoS derives the DSCP value by using the IP precedence value from the ingress packet and the IP-precedence-to-DSCP map. For non-IP packets that are tagged, QoS derives the DSCP value by using the received CoS value; for non-IP packets that are untagged, QoS derives the DSCP value by using the default port CoS value. In either case, the DSCP value is derived from the CoS-to-DSCP map. For more information, see the “Configuring the CoS-to-DSCP Map” section on page 26-59.
Step 5	set { ip dscp <i>new-dscp</i> ip precedence <i>new-precedence</i> }	Mark IP traffic by setting a new value in the packet: • For ip dscp <i>new-dscp</i>, enter a new DSCP value to be assigned to the classified traffic. The range is 0 to 63. • For ip precedence <i>new-precedence</i>, enter a new IP-precedence value to be assigned to the classified traffic. The range is 0 to 7.
Step 6	police <i>rate-bps burst-byte</i> [exceed-action { drop policed-dscp-transmit }]	Define a policer for the classified traffic. By default, no policer is defined. For information on the number of policers supported, see the “Standard QoS Configuration Guidelines” section on page 26-40. • For <i>rate-bps</i>, specify average traffic rate in bps. The range is 8000 to 1000000000. • For <i>burst-byte</i>, specify the normal burst size in bytes. The range is 8000 to 1000000. • (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 policed-dscp-transmit keywords to mark down the DSCP value (through the policed-DSCP map) and send the packet. For more information, see the “Configuring the Policed-DSCP Map” section on page 26-61.
Step 7	exit	Return to policy-map configuration mode.
Step 8	exit	Return to global configuration mode.

	Command	Purpose
Step 9	interface <i>interface-id</i>	Specify the port to attach to the policy map, and enter interface configuration mode. Valid interfaces include physical ports.
Step 10	service-policy input <i>policy-map-name</i>	Specify the ingress policy-map name, and apply it to a port. Only one policy map per port is supported.
Step 11	end	Return to privileged EXEC mode.
Step 12	show policy-map [<i>policy-map-name</i> [class <i>class-map-name</i>]]	Verify your entries.
Step 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, use the **no class** *class-name* policy-map configuration command. To return to the untrusted state, use the **no trust** policy-map class configuration command. To remove an assigned DSCP or IP precedence value, use the **no set {ip dscp new-dscp | ip precedence new-precedence}** policy-map class configuration command. To remove an existing policer, use the **no police rate-bps burst-byte [exceed-action {drop | policed-dscp-transmit}]** policy-map class configuration command. To remove the policy map and interface association, use the **no service-policy input** *policy-map-name* interface configuration command.

This example shows how to create an ingress policy-map and attach it to a port. 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 inbound packet is trusted. If the matched traffic exceeds an average traffic rate of 48000 bps and a normal burst size of 8000 bytes, its DSCP is marked down (based on the policed-DSCP map) 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)# trust dscp
Switch(config-pmap-c)# police 48000 8000 exceed-action policed-dscp-transmit
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/0/1
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 a port. 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. The second permit statement allows only Ethertype XNS-IDP traffic from the host with MAC address 0001.0000.0002 destined for the host with MAC address 0002.0000.0002.

```
Switch(config)# mac access-list extended maclist1
Switch(config-ext-mac)# permit 0001.0000.0001 0.0.0 0002.0000.0001 0.0.0
Switch(config-ext-mac)# permit 0001.0000.0002 0.0.0 0002.0000.0002 0.0.0 xns-idp
Switch(config-ext-mac)# exit
Switch(config)# mac access-list extended maclist2
Switch(config-ext-mac)# permit 0001.0000.0003 0.0.0 0002.0000.0003 0.0.0
Switch(config-ext-mac)# permit 0001.0000.0004 0.0.0 0002.0000.0004 0.0.0 aarp
Switch(config-ext-mac)# exit
```

```

Switch(config)# class-map macclass1
Switch(config-cmap)# match access-group maclist1
Switch(config-cmap)# exit
Switch(config)# policy-map macpolicy1
Switch(config-pmap)# class macclass1
Switch(config-pmap-c)# set ip dscp 63
Switch(config-pmap-c)# exit
Switch(config-pmap)# class macclass2 maclist2
Switch(config-pmap-c)# set ip dscp 45
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# mls qos trust cos
Switch(config-if)# service-policy input macpolicy1

```

Classifying, Policing, and Marking Ingress Traffic by Using Aggregate Policers

By using an aggregate policer, you can create a policer that is shared by multiple traffic classes within the same policy map. However, you cannot use the aggregate policer across different policy maps or ports.

Before beginning this procedure, make sure that you have created the class map to isolate traffic. For more information, see the [“Classifying Ingress Traffic by Using ACLs”](#) section on page 26-49 and the [“Classifying Ingress Traffic by Using Class Maps”](#) section on page 26-52.

Beginning in privileged EXEC mode, follow these steps to create an aggregate policer for inbound traffic:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos aggregate-policer <i>aggregate-policer-name rate-bps burst-byte</i> exceed-action {drop policed-dscp-transmit}	Define the policer parameters that can be applied to multiple traffic classes within the same policy map. By default, no aggregate policer is defined. For information on the number of policers supported, see the “Standard QoS Configuration Guidelines” section on page 26-40 - For <i>aggregate-policer-name</i>, specify the name of the aggregate policer. - For <i>rate-bps</i>, specify average traffic rate in bps. The range is 8000 to 1000000000. - For <i>burst-byte</i>, specify the normal burst size in bytes. The range is 8000 to 1000000. - Specify the action to take when the rates are exceeded. Use the exceed-action drop keywords to drop the packet. Use the exceed-action policed-dscp-transmit keywords to mark down the DSCP value (through the policed-DSCP map) and send the packet. For more information, see the “Configuring the Policed-DSCP Map” section on page 26-61.
Step 3	policy-map policy-map-name	Create a policy map by entering the policy-map name, and enter policy-map configuration mode. For more information, see the “Classifying, Policing, and Marking Ingress Traffic by Using Policy Maps” section on page 26-54.

	Command	Purpose
Step 4	class <i>class-name</i>	Specify the name of the class whose traffic policy you want to create or change, and enter policy-map class configuration mode. By default, no traffic classes are defined.
Step 5	police aggregate <i>aggregate-policer-name</i>	Apply an aggregate policer to multiple classes in the same policy map. For <i>aggregate-policer-name</i> , enter the name specified in Step 2.
Step 6	exit	Return to global configuration mode.
Step 7	interface <i>interface-id</i>	Specify the port to attach to the policy map, and enter interface configuration mode. Valid interfaces include physical ports.
Step 8	service-policy input <i>policy-map-name</i>	Specify the ingress policy-map name, and apply it to a port. Only one policy map per port is supported.
Step 9	end	Return to privileged EXEC mode.
Step 10	show mls qos aggregate-policer [<i>aggregate-policer-name</i>]	Verify your entries.
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the specified aggregate policer from a policy map, use the **no police aggregate** *aggregate-policer-name* policy-map class configuration command. To delete an aggregate policer and its parameters, use the **no mls qos aggregate-policer** *aggregate-policer-name* global configuration command.

This example shows how to create an aggregate policer and attach it to multiple classes within a policy map. In the configuration, the IP ACLs permit traffic from network 10.1.0.0 and from host 11.3.1.1. For traffic coming from network 10.1.0.0, the DSCP in the inbound packets is trusted. For traffic coming from host 11.3.1.1, the DSCP in the packet is changed to 56. The traffic rate from the 10.1.0.0 network and from host 11.3.1.1 is policed. If the traffic exceeds an average rate of 48000 bps and a normal burst size of 8000 bytes, its DSCP is marked down (based on the policed-DSCP map) and sent. The ingress policy-map is attached to a port.

```
Switch(config)# access-list 1 permit 10.1.0.0 0.0.255.255
Switch(config)# access-list 2 permit 11.3.1.1
Switch(config)# mls qos aggregate-police transmit1 48000 8000 exceed-action
policed-dscp-transmit
Switch(config)# class-map ipclass1
Switch(config-cmap)# match access-group 1
Switch(config-cmap)# exit
Switch(config)# class-map ipclass2
Switch(config-cmap)# match access-group 2
Switch(config-cmap)# exit
Switch(config)# policy-map aggflow1
Switch(config-pmap)# class ipclass1
Switch(config-pmap-c)# trust dscp
Switch(config-pmap-c)# police aggregate transmit1
Switch(config-pmap-c)# exit
Switch(config-pmap)# class ipclass2
Switch(config-pmap-c)# set ip dscp 56
Switch(config-pmap-c)# police aggregate transmit1
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
```

```
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# service-policy input aggflow1
Switch(config-if)# exit
```

Configuring DSCP Maps

These sections describe how to configure the DSCP maps:

- [Configuring the CoS-to-DSCP Map, page 26-59](#) (optional)
- [Configuring the IP-Precedence-to-DSCP Map, page 26-60](#) (optional)
- [Configuring the Policed-DSCP Map, page 26-61](#) (optional, unless the null settings in the map are not appropriate)
- [Configuring the DSCP-to-CoS Map, page 26-62](#) (optional)
- [Configuring the DSCP-to-DSCP-Mutation Map, page 26-63](#) (optional, unless the null settings in the map are not appropriate)

All the maps, except the DSCP-to-DSCP-mutation map, are globally defined and are applied to all ports.

Configuring the CoS-to-DSCP Map

You use the CoS-to-DSCP map to map CoS values in inbound packets to a DSCP value that QoS uses internally to represent the priority of the traffic.

[Table 26-12](#) shows the default CoS-to-DSCP map.

Table 26-12 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. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos map cos-dscp dscp1...dscp8	Modify the CoS-to-DSCP map. For <i>dscp1...dscp8</i> , enter eight DSCP values that correspond to CoS values 0 to 7. Separate each DSCP value with a space. The DSCP range is 0 to 63.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos maps cos-dscp	Verify your entries.
Step 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 cos-dscp** global configuration command.

This example shows how to modify and display the CoS-to-DSCP map:

```
Switch(config)# mls qos map cos-dscp 10 15 20 25 30 35 40 45
Switch(config)# end
Switch# show mls qos maps cos-dscp
```

```
Cos-dscp map:
  cos:   0  1  2  3  4  5  6  7
-----
  dscp:  10 15 20 25 30 35 40 45
```

Configuring the IP-Precedence-to-DSCP Map

You use the IP-precedence-to-DSCP map to map IP precedence values in inbound packets to a DSCP value that QoS uses internally to represent the priority of the traffic.

Table 26-13 shows the default IP-precedence-to-DSCP map:

Table 26-13 Default IP-Precedence-to-DSCP Map

IP precedence 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 IP-precedence-to-DSCP map. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos map ip-prec-dscp <i>dscp1...dscp8</i>	Modify the IP-precedence-to-DSCP map. For <i>dscp1...dscp8</i> , enter eight DSCP values that correspond to the IP precedence values 0 to 7. Separate each DSCP value with a space. The DSCP range is 0 to 63.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos maps ip-prec-dscp	Verify your entries.
Step 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 ip-prec-dscp** global configuration command.

This example shows how to modify and display the IP-precedence-to-DSCP map:

```
Switch(config)# mls qos map ip-prec-dscp 10 15 20 25 30 35 40 45
Switch(config)# end
Switch# show mls qos maps ip-prec-dscp
```

```
IpPrecedence-dscp map:
  ipprec:  0  1  2  3  4  5  6  7
-----
  dscp:   10 15 20 25 30 35 40 45
```


Configuring the Policed-DSCP Map

You use the policed-DSCP map to mark down a DSCP value to a new value as the result of an ingress policing and marking action.

The default policed-DSCP map is a null map, which maps an inbound DSCP value to the same DSCP value.

Beginning in privileged EXEC mode, follow these steps to modify the policed-DSCP map. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos map policed-dscp <i>dscp-list</i> to <i>mark-down-dscp</i>	Modify the policed-DSCP map. - For <i>dscp-list</i>, enter up to eight DSCP values separated by spaces. Then enter the to keyword. - For <i>mark-down-dscp</i>, enter the corresponding policed (marked down) DSCP value.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos maps policed-dscp	Verify your entries.
Step 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 policed-dscp** global configuration command.

This example shows how to map DSCP 50 to 57 to a marked-down DSCP value of 0:

```
Switch(config)# mls qos map policed-dscp 50 51 52 53 54 55 56 57 to 0
Switch(config)# end
Switch# show mls qos maps policed-dscp
Policed-dscp map:
  d1 :  d2 0  1  2  3  4  5  6  7  8  9
-----
  0 :    00 01 02 03 04 05 06 07 08 09
  1 :    10 11 12 13 14 15 16 17 18 19
  2 :    20 21 22 23 24 25 26 27 28 29
  3 :    30 31 32 33 34 35 36 37 38 39
  4 :    40 41 42 43 44 45 46 47 48 49
  5 :    00 00 00 00 00 00 00 00 58 59
  6 :    60 61 62 63
```



Note

In this policed-DSCP map, the marked-down DSCP values are shown in the body of the matrix. The d1 column specifies the most-significant digit of the original DSCP; the d2 row specifies the least-significant digit of the original DSCP. The intersection of the d1 and d2 values provides the marked-down value. For example, an original DSCP value of 53 corresponds to a marked-down DSCP value of 0.

Configuring the DSCP-to-CoS Map

You use the DSCP-to-CoS map to generate a CoS value, which is used to select one of the four queues in the egress queue-set.

Table 26-14 shows the default DSCP-to-CoS map.

Table 26-14 Default DSCP-to-CoS Map

DSCP value	0–7	8–15	16–23	24–31	32–39	40–47	48–55	56–63
CoS value	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. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 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 eight 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 DSCP range is 0 to 63; the CoS range is 0 to 7.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos maps dscp-to-cos	Verify your entries.
Step 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 dscp-cos** global configuration command.

This example shows how to map DSCP values 0, 8, 16, 24, 32, 40, 48, and 50 to CoS value 0 and to display the map:

```
Switch(config)# mls qos map dscp-cos 0 8 16 24 32 40 48 50 to 0
Switch(config)# end
Switch# show mls qos maps dscp-cos
Dscp-cos map:
  d1 :  d2 0  1  2  3  4  5  6  7  8  9
-----
  0 :    00 00 00 00 00 00 00 00 00 00 01
  1 :    01 01 01 01 01 01 00 02 02 02
  2 :    02 02 02 02 00 03 03 03 03 03
  3 :    03 03 00 04 04 04 04 04 04 04
  4 :    00 05 05 05 05 05 05 05 00 06
  5 :    00 06 06 06 06 06 07 07 07 07
  6 :    07 07 07 07
```



Note

In the above DSCP-to-CoS map, the CoS values are shown in the body of the matrix. The d1 column specifies the most-significant digit of the DSCP; the d2 row specifies the least-significant digit of the DSCP. The intersection of the d1 and d2 values provides the CoS value. For example, in the DSCP-to-CoS map, a DSCP value of 08 corresponds to a CoS value of 0.

Configuring the DSCP-to-DSCP-Mutation Map

If two QoS domains have different DSCP definitions, use the DSCP-to-DSCP-mutation map to translate one set of DSCP values to match the definition of another domain. You apply the DSCP-to-DSCP-mutation map to the receiving port (ingress mutation) at the boundary of a QoS administrative domain.

With ingress mutation, the new DSCP value overwrites the one in the packet, and QoS treats the packet with this new value. The switch sends the packet out the port with the new DSCP value.

You can configure multiple DSCP-to-DSCP-mutation maps and apply them to traffic received on a port. The default DSCP-to-DSCP-mutation map is a null map, which maps an inbound DSCP value to the same DSCP value.

Beginning in privileged EXEC mode, follow these steps to modify the DSCP-to-DSCP-mutation map. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos map dscp-mutation <i>dscp-mutation-name in-dscp to out-dscp</i>	Modify the DSCP-to-DSCP-mutation map. - For <i>dscp-mutation-name</i>, enter the mutation map name. You can create more than one map by specifying a new name. - For <i>in-dscp</i>, enter up to eight DSCP values separated by spaces. Then enter the to keyword. - For <i>out-dscp</i>, enter a single DSCP value. The DSCP range is 0 to 63.
Step 3	interface <i>interface-id</i>	Specify the port to which to attach the map, and enter interface configuration mode. Valid interfaces include physical ports.
Step 4	mls qos trust dscp	Configure the ingress port as a DSCP-trusted port. By default, the port is not trusted.
Step 5	mls qos dscp-mutation <i>dscp-mutation-name</i>	Apply the map to the specified ingress DSCP-trusted port. For <i>dscp-mutation-name</i> , enter the mutation map name specified in Step 2.
Step 6	end	Return to privileged EXEC mode.
Step 7	show mls qos maps dscp-mutation	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default map, use the **no mls qos dscp-mutation** *dscp-mutation-name* global configuration command.

This example shows how to define the DSCP-to-DSCP-mutation map. All the entries that are not explicitly configured are not modified (remains as specified in the null map).

```
Switch(config)# mls qos map dscp-mutation mutation1 1 2 3 4 5 6 7 to 0
Switch(config)# mls qos map dscp-mutation mutation1 8 9 10 11 12 13 to 10
Switch(config)# mls qos map dscp-mutation mutation1 20 21 22 to 20
Switch(config)# mls qos map dscp-mutation mutation1 30 31 32 33 34 to 30
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# mls qos trust dscp
```

```

Switch(config-if)# mls qos dscp-mutation mutation1
Switch(config-if)# end
Switch# show mls qos maps dscp-mutation mutation1
Dscp-dscp mutation map:
mutation1:
  d1 : d2 0  1  2  3  4  5  6  7  8  9
-----
  0 :   00 00 00 00 00 00 00 00 10 10
  1 :   10 10 10 10 14 15 16 17 18 19
  2 :   20 20 20 23 24 25 26 27 28 29
  3 :   30 30 30 30 30 35 36 37 38 39
  4 :   40 41 42 43 44 45 46 47 48 49
  5 :   50 51 52 53 54 55 56 57 58 59
  6 :   60 61 62 63

```

**Note**

In the above DSCP-to-DSCP-mutation map, the mutated values are shown in the body of the matrix. The d1 column specifies the most-significant digit of the original DSCP; the d2 row specifies the least-significant digit of the original DSCP. The intersection of the d1 and d2 values provides the mutated value. For example, a DSCP value of 12 corresponds to a mutated value of 10.

Configuring Ingress Queue Characteristics

Depending on the complexity of your network and your QoS solution, you might need to perform all of the tasks in the next sections. You will need to make decisions about these characteristics:

- Which packets are assigned (by DSCP or CoS value) to each queue?
- What drop percentage thresholds apply to each queue, and which CoS or DSCP values map to each threshold?
- How much of the available buffer space is allocated between the queues?
- How much of the available bandwidth is allocated between the queues?
- Is there traffic (such as voice) that should be given high priority?

These sections describe how to configure ingress queue characteristics:

- [Mapping DSCP or CoS Values to an Ingress Queue and Setting WTD Thresholds, page 26-65](#) (optional)
- [Allocating Buffer Space Between the Ingress Queues, page 26-66](#) (optional)
- [Allocating Bandwidth Between the Ingress Queues, page 26-67](#) (optional)
- [Configuring the Ingress Priority Queue, page 26-68](#) (optional)

Mapping DSCP or CoS Values to an Ingress Queue and Setting WTD Thresholds

You can prioritize inbound traffic by placing packets with particular DSCP or CoS values into certain queues and by adjusting the queue thresholds so that packets with lower priorities are dropped.

Beginning in privileged EXEC mode, follow these steps to map DSCP or CoS values to an ingress queue and to set WTD thresholds. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos srr-queue input dscp-map queue <i>queue-id</i> threshold <i>threshold-id</i> dscp1...dscp8 or mls qos srr-queue input cos-map queue <i>queue-id</i> threshold <i>threshold-id</i> cos1...cos8	Map DSCP or CoS values to an ingress queue and to a threshold ID. By default, DSCP values 0–39 and 48–63 are mapped to queue 1 and threshold 1. DSCP values 40–47 are mapped to queue 2 and threshold 1. By default, CoS values 0–4, 6, and 7 are mapped to queue 1 and threshold 1. CoS value 5 is mapped to queue 2 and threshold 1. - For <i>queue-id</i>, the range is 1 to 2. - For <i>threshold-id</i>, the range is 1 to 3. The drop-threshold percentage for threshold 3 is predefined. It is set to the queue-full state. - For <i>dscp1...dscp8</i>, enter up to eight values, and separate each value with a space. The range is 0 to 63. - For <i>cos1...cos8</i>, enter up to eight values, and separate each value with a space. The range is 0 to 7.
Step 3	mls qos srr-queue input threshold <i>queue-id</i> <i>threshold-percentage1</i> <i>threshold-percentage2</i>	Assign the two WTD threshold percentages for (threshold 1 and 2) to an ingress queue. The default, both thresholds are set to 100 percent. - For <i>queue-id</i>, the range is 1 to 2. - For <i>threshold-percentage1 threshold-percentage2</i>, the range is 1 to 100. Separate each value with a space. Each threshold value is a percentage of the total number of queue descriptors allocated for the queue.
Step 4	end	Return to privileged EXEC mode.
Step 5	show mls qos maps	Verify your entries. The DSCP input queue threshold map appears as a matrix. The d1 column specifies the most-significant digit of the DSCP number; the d2 row specifies the least-significant digit in the DSCP number. The intersection of the d1 and the d2 values provides the queue ID and threshold ID; for example, queue 2 and threshold 1 (02-01). The CoS input queue threshold map shows the CoS value in the top row and the corresponding queue ID and threshold ID in the second row; for example, queue 2 and threshold 2 (2-2).
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default CoS input queue threshold map or the default DSCP input queue threshold map, use the **no mls qos srr-queue input cos-map** or the **no mls qos srr-queue input dscp-map** global configuration command. To return to the default WTD threshold percentages, use the **no mls qos srr-queue input threshold *queue-id*** global configuration command.

This example shows how to map DSCP values 0 to 6 to ingress queue 1 and to threshold 1 with a drop threshold of 50 percent. It maps DSCP values 20 to 26 to ingress queue 1 and to threshold 2 with a drop threshold of 70 percent.

```
Switch(config)# mls qos srr-queue input dscp-map queue 1 threshold 1 0 1 2 3 4 5 6
Switch(config)# mls qos srr-queue input dscp-map queue 1 threshold 2 20 21 22 23 24 25 26
Switch(config)# mls qos srr-queue input threshold 1 50 70
```

In this example, the DSCP values (0 to 6) are assigned the WTD threshold of 50 percent and will be dropped sooner than the DSCP values (20 to 26) assigned to the WTD threshold of 70 percent.

Allocating Buffer Space Between the Ingress Queues

You define the ratio (allocate the amount of space) with which to divide the ingress buffers between the two queues. The buffer and the bandwidth allocation control how much data can be buffered before packets are dropped.

Beginning in privileged EXEC mode, follow these steps to allocate the buffers between the ingress queues. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos srr-queue input buffers <i>percentage1 percentage2</i>	Allocate the buffers between the ingress queues By default 90 percent of the buffers are allocated to queue 1, and 10 percent of the buffers are allocated to queue 2. For <i>percentage1 percentage2</i> , the range is 0 to 100. Separate each value with a space. You should allocate the buffers so that the queues can handle any inbound bursty traffic.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos interface buffer or show mls qos input-queue	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no mls qos srr-queue input buffers** global configuration command.

This example shows how to allocate 60 percent of the buffer space to ingress queue 1 and 40 percent of the buffer space to ingress queue 2:

```
Switch(config)# mls qos srr-queue input buffers 60 40
```

Allocating Bandwidth Between the Ingress Queues

You need to specify how much of the available bandwidth is allocated between the ingress queues. The ratio of the weights is the ratio of the frequency in which the SRR scheduler sends packets from each queue to the internal ring. The bandwidth and the buffer allocation control how much data can be buffered before packets are dropped. On ingress queues, SRR operates only in shared mode.

Beginning in privileged EXEC mode, follow these steps to allocate bandwidth between the ingress queues. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos srr-queue input bandwidth <i>weight1 weight2</i>	Assign shared round robin weights to the ingress queues. The default setting for <i>weight1</i> and <i>weight2</i> is 4 (1/2 of the bandwidth is equally shared between the two queues). For <i>weight1</i> and <i>weight2</i> , the range is 1 to 100. Separate each value with a space. SRR services the priority queue for its configured weight as specified by the bandwidth keyword in the mls qos srr-queue input priority-queue queue-id bandwidth weight global configuration command. Then, SRR shares the remaining bandwidth with both ingress queues and services them as specified by the weights configured with the mls qos srr-queue input bandwidth weight1 weight2 global configuration command. For more information, see the “Configuring the Ingress Priority Queue” section on page 26-68 .
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos interface queueing or show mls qos input-queue	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no mls qos srr-queue input bandwidth** global configuration command.

This example shows how to assign the ingress bandwidth to the queues. Priority queueing is disabled, and the shared bandwidth ratio allocated to queue 1 is 25/(25+75) and to queue 2 is 75/(25+75).

```
Switch(config)# mls qos srr-queue input priority-queue 2 bandwidth 0
Switch(config)# mls qos srr-queue input bandwidth 25 75
```

Configuring the Ingress Priority Queue

You should use the ingress priority queue only for traffic that needs to be expedited (for example, voice traffic, which needs minimum delay and jitter).

The priority queue is guaranteed part of the bandwidth to reduce the delay and jitter under heavy network traffic on an oversubscribed ring (when there is more traffic than the backplane can carry, and the queues are full and dropping frames).

SRR services the priority queue for its configured weight as specified by the **bandwidth** keyword in the **mls qos srr-queue input priority-queue queue-id bandwidth weight** global configuration command. Then, SRR shares the remaining bandwidth with both ingress queues and services them as specified by the weights configured with the **mls qos srr-queue input bandwidth weight1 weight2** global configuration command.

Beginning in privileged EXEC mode, follow these steps to configure the ingress priority queue. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos srr-queue input priority-queue queue-id bandwidth weight	Assign a queue as the priority queue and guarantee bandwidth on the internal ring if the ring is congested. By default, the priority queue is queue 2, and 10 percent of the bandwidth is allocated to it. - For <i>queue-id</i>, the range is 1 to 2. - For bandwidth weight, assign the bandwidth percentage of the internal ring. The range is 0 to 40. The amount of bandwidth that can be guaranteed is restricted because a large value affects the entire ring and can degrade performance.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos interface queueing or show mls qos input-queue	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no mls qos srr-queue input priority-queue queue-id** global configuration command. To disable priority queueing, set the bandwidth weight to 0, for example, **mls qos srr-queue input priority-queue queue-id bandwidth 0**.

This example shows how to assign the ingress bandwidths to the queues. Queue 1 is the priority queue with 10 percent of the bandwidth allocated to it. The bandwidth ratios allocated to queues 1 and 2 is 4/(4+4). SRR services queue 1 (the priority queue) first for its configured 10 percent bandwidth. Then SRR equally shares the remaining 90 percent of the bandwidth between queues 1 and 2 by allocating 45 percent to each queue.

```
Switch(config)# mls qos srr-queue input priority-queue 1 bandwidth 10
Switch(config)# mls qos srr-queue input bandwidth 4 4
```


Configuring Egress Queue-Set Characteristics

Depending on the complexity of your network and your QoS solution, you might need to perform all of the tasks in the next sections. You will need to make decisions about these characteristics:

- Which packets are mapped by DSCP or CoS value to each queue and threshold ID?
- What drop percentage thresholds apply to the queue-set (four egress queues per port), and how much reserved and maximum memory is needed for the traffic type?
- How much of the fixed buffer space is allocated to the queue-set?
- Does the bandwidth of the port need to be rate limited?
- How often should the egress queue-set be serviced and which technique (shaped, shared, or both) should be used?

These sections describe how to configure egress queue-set characteristics:

- [Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set, page 26-69](#) (optional)
- [Mapping DSCP or CoS Values to an Egress Queue-Set and to a Threshold ID, page 26-71](#) (optional)
- [Configuring SRR Shaped Weights on an Egress Queue-Set, page 26-72](#) (optional)
- [Configuring SRR Shared Weights on an Egress Queue-Set, page 26-73](#) (optional)
- [Limiting the Egress Bandwidth on a Queue-Set, page 26-74](#) (optional)

Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set

You can guarantee the availability of buffers, set WTD thresholds, and configure the maximum memory allocation for a queue-set by using the **mls qos queue-set output *qset-id* threshold *queue-id* drop-threshold1 drop-threshold2 reserved-threshold maximum-threshold** global configuration command.

Each threshold value is a percentage of the queue's allocated memory, which you specify by using the **mls qos queue-set output *qset-id* buffers *allocation1* ... *allocation4*** global configuration command. The queues use WTD to support distinct drop percentages for different traffic classes.



Note

The egress queue-set default settings are suitable for most situations. You should change them only when you have a thorough understanding of the queues and only if these settings do not meet your QoS solution.

Beginning in privileged EXEC mode, follow these steps to configure the memory allocation and drop thresholds for a queue-set. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos queue-set output <i>qset-id</i> buffers <i>allocation1 ... allocation4</i>	Allocate buffers to a queue-set. By default, all allocation values are equally mapped among the four queues (25, 25, 25, 25). Each queue has 1/4 of the buffer space. - For <i>qset-id</i>, enter the ID of the queue-set. The range is 1 to 2. Each port belongs to a queue-set, which defines all the characteristics of the four egress queues per port. - For <i>allocation1 ... allocation4</i>, specify four percentages, one for each queue in the queue-set. The range is 0 to 100. Separate each value with a space. Allocate buffers according to the importance of the traffic. For example, give a larger percentage of the buffer to the queue with the highest-priority traffic.
Step 3	mls qos queue-set output <i>qset-id</i> threshold <i>queue-id</i> <i>drop-threshold1</i> <i>drop-threshold2</i> <i>reserved-threshold</i> <i>maximum-threshold</i>	Configure the WTD thresholds, guarantee the availability of buffers, and configure the maximum memory allocation for the queue-set (four egress queues per port). By default, the WTD thresholds for queues 1, 3, and 4 are set to 100 percent. The thresholds for queue 2 are set to 50 percent. The reserved thresholds for queues 1, 3, and 4 are set to 50 percent. The reserved threshold for queue 2 is set to 100 percent. The maximum thresholds for all queues are set to 400 percent. - For <i>qset-id</i>, enter the ID of the queue-set specified in Step 2. The range is 1 to 2. - For <i>queue-id</i>, enter the specific queue in the queue-set on which the command is performed. The range is 1 to 4. - For <i>drop-threshold1 drop-threshold2</i>, specify the two WTD thresholds expressed as a percentage of the queue's allocated memory. The range is 1 to 400 percent. - For <i>reserved-threshold</i>, enter the amount of memory to be guaranteed (reserved) for the queue expressed as a percentage of the allocated memory. The range is 1 to 100 percent. - For <i>maximum-threshold</i>, enable a full queue to obtain more buffers than are reserved for it. This is the maximum memory the queue can have before the packets are dropped if the common pool is not empty. The range is 1 to 400 percent.
Step 4	interface <i>interface-id</i>	Specify the port, and enter interface configuration mode.
Step 5	queue-set <i>qset-id</i>	Map the port to a queue-set. For <i>qset-id</i>, enter the ID of the queue-set specified in Step 2. The range is 1 to 2. The default is 1.
Step 6	end	Return to privileged EXEC mode.
Step 7	show mls qos interface [<i>interface-id</i>] buffers	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no mls qos queue-set output *qset-id* buffers** global configuration command. To return to the default WTD threshold percentages, use the **no mls qos queue-set output *qset-id* threshold [*queue-id*]** global configuration command.

This example shows how to map a port to queue-set 2. It allocates 40 percent of the buffer space to egress queue 1 and 20 percent each to egress queues 2, 3, and 4. It configures the drop thresholds for queue 2 to 40 and 60 percent of the allocated memory, guarantees (reserves) 100 percent of the allocated memory, and configures 200 percent as the maximum memory that this queue can have before packets are dropped.

```
Switch(config)# mls qos queue-set output 2 buffers 40 20 20 20
Switch(config)# mls qos queue-set output 2 threshold 2 40 60 100 200
Switch(config)# interface fastethernet1/0/1
Switch(config-if)# queue-set 2
```

Mapping DSCP or CoS Values to an Egress Queue-Set and to a Threshold ID

You can prioritize traffic by placing packets with particular DSCP or CoS values into certain queues and by adjusting the queue thresholds so that packets with lower priorities are dropped.



Note

The egress queue-set default settings are suitable for most situations. You should change them only when you have a thorough understanding of the queues and only if these settings do not meet your QoS solution.

Beginning in privileged EXEC mode, follow these steps to map DSCP or CoS values to an egress queue and to a threshold ID. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos srr-queue output dscp-map queue <i>queue-id</i> threshold <i>threshold-id</i> <i>dscp1...dscp8</i> or mls qos srr-queue output cos-map queue <i>queue-id</i> threshold <i>threshold-id</i> <i>cos1...cos8</i>	Map DSCP or CoS values to an egress queue and to a threshold ID for a port. By default, DSCP values 0–15 are mapped to queue 2 and threshold 1. DSCP values 16–31 are mapped to queue 3 and threshold 1. DSCP values 32–39 and 48–63 are mapped to queue 4 and threshold 1. DSCP values 40–47 are mapped to queue 1 and threshold 1. By default, CoS values 0 and 1 are mapped to queue 2 and threshold 1. CoS values 2 and 3 are mapped to queue 3 and threshold 1. CoS values 4, 6, and 7 are mapped to queue 4 and threshold 1. CoS value 5 is mapped to queue 1 and threshold 1. - For <i>queue-id</i>, the range is 1 to 4. - For <i>threshold-id</i>, the range is 1 to 3. The drop-threshold percentage for threshold 3 is predefined. It is set to the queue-full state. - For <i>dscp1...dscp8</i>, enter up to eight values, and separate each value with a space. The range is 0 to 63. - For <i>cos1...cos8</i>, enter up to eight values, and separate each value with a space. The range is 0 to 7.
Step 3	end	Return to privileged EXEC mode.

	Command	Purpose
Step 4	show mls qos maps	Verify your entries. The DSCP output queue threshold map appears as a matrix. The d1 column specifies the most-significant digit of the DSCP number; the d2 row specifies the least-significant digit in the DSCP number. The intersection of the d1 and the d2 values provides the queue ID and threshold ID; for example, queue 2 and threshold 1 (02-01). The CoS output queue threshold map shows the CoS value in the top row and the corresponding queue ID and threshold ID in the second row; for example, queue 2 and threshold 2 (2-2).
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default DSCP output queue threshold map or the default CoS output queue threshold map, use the **no mls qos srr-queue output dscp-map** or the **no mls qos srr-queue output cos-map** global configuration command.

This example shows how to map DSCP values 10 and 11 to egress queue 1 and to threshold 2:

```
Switch(config)# mls qos srr-queue output dscp-map queue 1 threshold 2 10 11
```

Configuring SRR Shaped Weights on an Egress Queue-Set

You can specify how much of the available bandwidth is allocated to each queue in the queue-set. The ratio of the weights is the ratio of frequency in which the SRR scheduler sends packets from a standard port.



Note

SRR shaping is not supported on the ES ports; however, you can configure average-rate shaping. For more information, see the [“Configuring Shaping” section on page 26-99](#).

You can configure the egress queues for shaped weights, shared weights, or both. Use shaping to smooth bursty traffic or to provide a smoother output over time. For conceptual information, see the [“SRR Shaping and Sharing” section on page 26-14](#). For configuration information, see the [“Configuring SRR Shared Weights on an Egress Queue-Set” section on page 26-73](#).



Note

The egress queue-set default settings are suitable for most situations. You should change them only when you have a thorough understanding of the queues and only if these settings do not meet your QoS solution.

Beginning in privileged EXEC mode, follow these steps to assign the shaped weights and to enable bandwidth shaping on a standard port mapped to the four egress queues. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify a standard port, and enter interface configuration mode.

	Command	Purpose
Step 3	srr-queue bandwidth shape <i>weight1 weight2 weight3 weight4</i>	Assign SRR weights to the egress queues. This command is not supported on an ES port. By default, weight1 is set to 25; weight2, weight3, and weight4 are set to 0, and these queues are in shared mode. For <i>weight1 weight2 weight3 weight4</i>, enter the weights to control the percentage of the port that is shaped. The inverse ratio ($1/\text{weight}$) controls the shaping bandwidth for this queue. Separate each value with a space. The range is 0 to 65535. If you configure a weight of 0, the corresponding queue operates in shared mode. The weight specified with the srr-queue bandwidth shape command is ignored, and the weights specified with the srr-queue bandwidth share interface configuration command for a queue come into effect. When configuring queues in the same queue-set for both shaping and sharing, make sure that you configure the lowest number queue for shaping. The shaped mode overrides the shared mode.
Step 4	end	Return to privileged EXEC mode.
Step 5	show mls qos interface <i>interface-id</i> queueing	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no srr-queue bandwidth shape** interface configuration command.

This example shows how to configure bandwidth shaping on queue 1. Because the weight ratios for queues 2, 3, and 4 are set to 0, these queues operate in shared mode. The bandwidth weight for queue 1 is 1/8, which is 12.5 percent.

```
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# srr-queue bandwidth shape 8 0 0 0
```

Configuring SRR Shared Weights on an Egress Queue-Set

In shared mode, the queues in the queue-set share the bandwidth among them according to the configured weights. The bandwidth is guaranteed at this level but not limited to it. For example, if a queue empties and does not require a share of the link, the remaining queues can expand into the unused bandwidth and share it among them. With sharing, the ratio of the weights controls the frequency of dequeuing; the absolute values are meaningless.



Note

The egress queue-set default settings are suitable for most situations. You should change them only when you have a thorough understanding of the queues and only if these settings do not meet your QoS solution.

Beginning in privileged EXEC mode, follow these steps to assign the shared weights and to enable bandwidth sharing on a port mapped to the four egress queues. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify a port, and enter interface configuration mode.
Step 3	srr-queue bandwidth share <i>weight1 weight2 weight3 weight4</i>	Assign SRR weights to the egress queues. By default, all four weights are 25 (1/4 of the bandwidth is allocated to each queue). For <i>weight1 weight2 weight3 weight4</i> , enter the weights to control the ratio of the frequency in which the SRR scheduler sends packets. Separate each value with a space. The range is 1 to 255.
Step 4	end	Return to privileged EXEC mode.
Step 5	show mls qos interface <i>interface-id</i> queueing	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no srr-queue bandwidth share** interface configuration command.

This example shows how to configure the weight ratio of the SRR scheduler running on an egress port. Four queues are used, and the bandwidth ratio allocated for each queue in shared mode is $1/(1+2+3+4)$, $2/(1+2+3+4)$, $3/(1+2+3+4)$, and $4/(1+2+3+4)$, which is 10 percent, 20 percent, 30 percent, and 40 percent for queues 1, 2, 3, and 4. This means that queue 4 has four times the bandwidth of queue 1, twice the bandwidth of queue 2, and one-and-a-third times the bandwidth of queue 3.

```
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# srr-queue bandwidth share 1 2 3 4
```

Limiting the Egress Bandwidth on a Queue-Set

You can limit the egress bandwidth on a standard port mapped to a queue-set. For example, if a customer pays only for a small percentage of a high-speed link, you can limit the bandwidth to that amount.



Note

The egress queue-set default settings are suitable for most situations. You should change them only when you have a thorough understanding of the queues and only if these settings do not meet your QoS solution.

Beginning in privileged EXEC mode, follow these steps to limit the egress bandwidth on a standard port. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify a standard port to be rate-limited, and enter interface configuration mode.
Step 3	srr-queue bandwidth limit <i>weight1</i>	Specify the percentage of the port speed to which the port should be limited. The range is 10 to 90. This command is not supported on an ES port. By default, the port is not rate limited and is set to 100 percent.

	Command	Purpose
Step 4	end	Return to privileged EXEC mode.
Step 5	show mls qos interface <i>[interface-id]</i> queueing	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no srr-queue bandwidth limit** interface configuration command.

This example shows how to limit the bandwidth on a standard port to 80 percent:

```
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# srr-queue bandwidth limit 80
```

When you configure this command to 80 percent, the port is idle 20 percent of the time. The line rate drops to 80 percent of the connected speed, which is 800 Mbps. These values are not exact because the hardware adjusts the line rate in increments of six.

Displaying Standard QoS Information

To display standard QoS information, use one or more of the privileged EXEC commands in [Table 26-15](#):

Table 26-15 Commands for Displaying Standard 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 mls qos	Display global QoS configuration information.
show mls qos aggregate-policer <i>[aggregate-policer-name]</i>	Display the aggregate policer configuration.
show mls qos input-queue	Display QoS settings for the ingress queues.
show mls qos interface <i>[interface-id]</i> [buffers policers queueing statistics]	Display QoS information at the port level, including the buffer allocation, which ports have policers, the queueing strategy, and the ingress and egress statistics.
show mls qos maps [cos-dscp cos-input-q cos-output-q dscp-cos dscp-input-q dscp-mutation <i>dscp-mutation-name</i> dscp-output-q ip-prec-dscp policed-dscp]	Display QoS mapping information.
show mls qos queue-set <i>[qset-id]</i>	Display the egress queue-set settings.
show policy-map <i>[policy-map-name [class class-map-name]]</i>	Display QoS policy-maps, which define the traffic policy for a traffic class.
show policy-map interface <i>interface-id</i> [input]	Display the ingress policy-map name applied to the specified port.

Configuring Hierarchical QoS

You can configure hierarchical QoS (traffic classification, CBWFQ, LLQ, shaping, and two-rate, three-color policing) and apply it to outbound traffic on an ES port.

Before configuring hierarchical 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.

These sections describe how to configure hierarchical QoS on your switch. Read these sections if traffic is flowing from a standard or an ES port to an ES port:

- [Default Hierarchical QoS Configuration, page 26-76](#)
- [Hierarchical QoS Configuration Guidelines, page 26-76](#)
- [Configuring an Egress Hierarchical QoS Policy, page 26-77](#) (required)

Default Hierarchical QoS Configuration

QoS is disabled.

No traffic classes, class maps, policy maps, or policers are defined.

LLQ is disabled.

CBWFQ is disabled.

Tail drop is enabled.

WRED is disabled.

Average-rate traffic shaping is disabled.

Hierarchical QoS Configuration Guidelines

Before beginning the hierarchical QoS configuration, you should be aware of this information:

- QoS must be enabled with the **mls qos** global configuration command for any hierarchical configuration to take effect. When enabled, QoS uses the default settings described in the “[Default Standard QoS Configuration](#)” section on page 26-38 and the “[Default Hierarchical QoS Configuration](#)” section on page 26-76. For detailed steps, see the “[Enabling QoS Globally](#)” section on page 26-42.
- Because the switch does not support attaching a service policy to a logical interface (such as an EtherChannel), 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.
- Class maps that contain ACLs are not supported in an egress policy attached to an ES port. You cannot use the **match access-group acl-index-or-name** class-map configuration command.

- To configure a class map that matches an 802.1Q tunneling pair (instead of matching a single VLAN), you must configure the **class-map** global configuration command with the **match-all** keyword. You must enter the **match vlan** command before the **match vlan inner** command.
- You cannot have **match vlan {vlan-id | inner vlan-id}** and **match {cos cos-list | ip dscp dscp-list | ip precedence ip-precedence-list | mpls experimental exp-list}** class-map configuration commands in the same class map.
- You cannot mix VLAN-level and class-level matches within a class map.
- When port trust policies are used with 802.1Q tunneling, all ports sharing the same tunnel VLAN must be configured with the same trust policy.
- Only one policy map per port is supported. You can attach one ingress service-policy per port and one egress service-policy per ES port.
- On an ES port, there is no limit to the number of egress policers.
- You can define a class policy to use either tail drop through the **queue-limit** policy-map class configuration command or to use WRED packet drop through the **random-detect** policy-map class configuration command. You cannot use the **queue-limit** and **random-detect** commands in the same class policy, but they can be used in two class policies in the same policy map.
- You cannot use **bandwidth**, **queue-limit**, **random-detect**, and **shape** policy-map class configuration commands with the **priority** policy-map class configuration command in the same class within the same policy map. However, you can use these commands in different classes in the same policy map. Within a policy map, you can give priority status to only one class.
- You must configure the **bandwidth** or the **shape** policy-map class configuration command before you configure either the **queue-limit** or the **random-detect** policy-map class configuration command in a class policy. You must configure the **bandwidth** or the **shape** command in the same policy map as the **queue-limit** or the **random-detect** command if the policy is not using the default traffic class. If the policy is using the default traffic class, you do not need to specify the **bandwidth** and **shape** commands in the policy map.
- A policy map can have all the class bandwidths specified in either kbps or in percentages, but not a mix of both. You cannot specify bandwidth in kbps in a child policy (configured through the **service-policy** policy-map class configuration command) and then specify bandwidth as a percentage in the parent policy.

Configuring an Egress Hierarchical QoS Policy

These sections describe how to configure an egress hierarchical QoS policy to create a service policy that is attached to an ES port.

For background information, see the [“Hierarchical Levels”](#) section on page 26-20, the [“Egress Classification Based on Traffic Classes and Traffic Policies”](#) section on page 26-23, the [“Egress Policing and Marking”](#) section on page 26-24, and the [“Queueing and Scheduling of Hierarchical Queues”](#) section on page 26-26. For configuration guidelines, see the [“Hierarchical QoS Configuration Guidelines”](#) section on page 26-76.

Depending on your network configuration and QoS solution, you must perform one or more of these tasks if you want classify, policy, mark, queue, or schedule outbound traffic:

- [Classifying Egress Traffic by Using Class Maps, page 26-78](#) (required)
- [Configuring an Egress Two-Rate Traffic Policer, page 26-80](#) (optional)
- [Configuring Class-Based Packet Marking in an Egress Traffic Policy, page 26-84](#) (optional)

- [Configuring CBWFQ and Tail Drop, page 26-86](#) (optional)
- [Configuring CBWFQ and DSCP-Based WRED, page 26-89](#) (optional)
- [Configuring CBWFQ and IP Precedence-Based WRED, page 26-93](#) (optional)
- [Enabling LLQ, page 26-97](#) (optional)
- [Configuring Shaping, page 26-99](#) (optional)

Classifying Egress Traffic by Using Class Maps

You use the **class-map** global configuration command to create a class map for matching packets to the class whose name you specify. The class map isolates a specific outbound traffic flow (class) from all other traffic by defining the criteria to use to match against a specific flow. The match criterion is defined with a match statement entered within the class-map configuration mode. Packets are checked against the match criteria configured for a class map. If a packet matches the specified criteria, the packet is considered a member of the class and is forwarded according to the QoS specifications set in the traffic policy.

Beginning in privileged EXEC mode, follow these steps to create a class-level class-map and to define the match criterion to classify outbound traffic. This procedure is required. The examples that follow the procedure show how to create a class-level and a VLAN-level class-map.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	class-map [match-all match-any] <i>class-map-name</i>	Create a class map, and enter class-map configuration mode. By default, no class maps are defined. • (Optional) Use the match-all keyword to perform a logical-AND of all matching statements under this class map. All criteria in the class map must be matched. • (Optional) Use the match-any keyword to perform a logical-OR of all matching statements under this class map. One or more criteria must be matched. • For <i>class-map-name</i>, specify the name of the class map. If neither the match-all nor the match-any keyword is specified, the default is match-all. You must use the match-all keyword if you are matching an 802.1Q tunneling pair (instead of matching a single VLAN).

	Command	Purpose
Step 3	match { cos <i>cos-list</i> ip dscp <i>dscp-list</i> ip precedence <i>ip-precedence-list</i> mpls experimental <i>exp-list</i> vlan { <i>vlan-id</i> inner <i>vlan-id</i> }}	Define the match criterion to classify traffic. By default, no match criterion is defined. - For cos <i>cos-list</i>, specify up to four Layer 2 CoS values to match against the packet. Separate each value with a space. The range is 0 to 7. - For ip dscp <i>dscp-list</i>, specify up to eight IP DSCP values to match against the packet. Separate each value with a space. The range is 0 to 63. - For ip precedence <i>ip-precedence-list</i>, specify up to eight IP-precedence values to match against the packet. Separate each value with a space. The range is 0 to 7. - For mpls experimental <i>exp-list</i>, specify up to eight MPLS EXP values to match against the packet. Separate each value with a space. The range is 0 to 7. - For vlan {<i>vlan-id</i> inner <i>vlan-id</i>}, specify packet match based on the VLAN ID. - For vlan <i>vlan-id</i>, match a packet based on the VLAN ID, or match a packet based on the outer VLAN ID if an 802.1Q tunnel is configured. You can specify a single VLAN identified by a VLAN number or a range of VLANs separated by a hyphen. The range is 1 to 4094. - For vlan inner <i>vlan-id</i>, match a packet based on the inner VLAN ID of an 802.1Q tunnel. You can specify a single VLAN identified by a VLAN number or a range of VLANs separated by a hyphen. The range is 1 to 4094. When matching an 802.1Q tunneling pair (instead of matching a single VLAN), you must enter the match vlan command before the match vlan inner command. Note You cannot mix VLAN-level and class-level matches within a class map.
Step 4	end	Return to privileged EXEC mode.
Step 5	show class-map	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete an existing class map, use the **no class-map** [**match-all** | **match-any**] *class-map-name* global configuration command. To remove a match criterion, use the **no match** {**cos** *cos-list* | **ip dscp** *dscp-list* | **ip precedence** *ip-precedence-list* | **mpls experimental** *exp-list* | **vlan** {*vlan-id* | **inner** *vlan-id*}} class-map configuration command.

This example shows how to create a class-level class-map called *class3*, which matches traffic with IP-precedence values of 5, 6, and 7:

```
Switch(config)# class-map class3
Switch(config-cmap)# match ip precedence 5 6 7
```

This example shows how to create a VLAN-level class-map for 802.1Q tunneling called *dot1q*, which matches all traffic with an outer VLAN ID of 5 and an inner VLAN ID of 3 to 8:

```
Switch(config)# class-map match-all dot1q
Switch(config-cmap)# match vlan 5
Switch(config-cmap)# match vlan inner 3 - 8
```

This example shows how to create a VLAN-level class-map called *vlan203*, which matches traffic in VLAN 203:

```
Switch(config)# class-map match-all vlan203
Switch(config-cmap)# match vlan 203
```

Configuring an Egress Two-Rate Traffic Policer

You can configure a two-rate traffic policer within a policy map at the class level, at the VLAN level, and at the physical level by using the **police cir** or the **police cir percent** policy-map class configuration command. At the physical level of the hierarchy, you can police only the class-default class in an egress policy attached to an ES port.

The policer limits the transmission rate of a traffic class and marks actions (conform, exceed, and violate) for each packet. Within the conform, exceed, and violate categories, you decide packet treatments. In the most common configurations, you configure packets that conform to be sent, packets that exceed to be sent with a decreased priority, and packets that violate to be dropped. You can decrease the priority of the CoS, the DSCP, the IP precedence, or the MPLS EXP bits in the packet. You configure the policer by using the **police cir** policy-map class configuration command.

You also can use the **police cir percent** policy-map class configuration command to configure an egress traffic policer that uses two rates, the CIR and the PIR. The switch calculates the CIR and PIR based on a percentage of the maximum amount of bandwidth assigned to the parent class. The maximum bandwidth is controlled by the **shape** policy-map class configuration command if it is configured. Otherwise, the maximum bandwidth is the physical port bandwidth (1 Gbps). For more information about the interaction between the **police cir percent** command and the **shape** command, refer to the **police cir percent** command in the command reference for this release.

Before beginning this procedure, make sure that you have created the class map to isolate traffic. For more information, see the [“Classifying Egress Traffic by Using Class Maps”](#) section on page 26-78.

Beginning in privileged EXEC mode, follow these steps to configure a class-level, two-rate traffic policer in a service policy for outbound traffic. This procedure is optional. The examples that follow the procedure show how to configure a class-level, a VLAN-level, and a physical-level policer.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	policy-map <i>policy-map-name</i>	Create a policy map by entering the policy-map name, and enter policy-map configuration mode. By default, no policy maps are defined.
Step 3	class <i>class-name</i>	Specify the name of the class whose traffic policy you want to create or change, and enter policy-map class configuration mode. By default, no traffic classes are defined.

	Command	Purpose
Step 4	police cir <i>cir</i> [bc <i>conform-burst</i>] pir <i>pir</i> [be <i>peak-burst</i>] [conform-action <i>action</i> [exceed-action <i>action</i> [violate-action <i>action</i>]]] or police cir percent <i>percent</i> [bc <i>conform-burst ms</i>] pir percent <i>percent</i> [be <i>peak-burst ms</i>] [conform-action <i>action</i> [exceed-action <i>action</i> [violate-action <i>action</i>]]] Note For the syntax description of the police cir percent command, refer to the command reference for this release.	Configure an egress traffic policer that uses two rates, the CIR and the peak information rate PIR. By default, no policer is defined. - For cir <i>cir</i>, specify the CIR at which the first token bucket is updated. The range is 64000 to 990000000 bps. (Optional) For bc <i>conform-burst</i>, specify the conform burst size used by the first token bucket for policing. The range is 1536 to 16776960 bytes. The default is 8192. - For pir <i>pir</i>, specify the PIR at which the second token bucket for policing is updated. The range is 64000 to 990000000 bps. (Optional) For be <i>peak-burst</i>, specify the peak burst size used by the second token bucket. The range is 1536 to 16776960 bytes. The default is 8192. (Optional) For conform-action, specify the action to perform on packets that conform to the CIR and PIR. (Optional) For exceed-action, specify the action to perform on packets that conform to the PIR but not the CIR. (Optional) For violate-action, specify the action to perform on packets that exceed the PIR. Note If you do not specify an action, the default conform-action is transmit, the default exceed-action is drop, and the default violate-action is drop. (Optional) For <i>action</i>, specify the action to perform on packets: drop—drop the packet. set-cos-transmit <i>new-cos</i>—set the CoS value to a new value, and send the packet. The range is 0 to 7. set-dscp-transmit <i>new-dscp</i>—set the IP DSCP value to a new value, and send the packet. The range is 0 to 63. You also can enter a mnemonic name for a commonly used value. set-mpls-exp-transmit <i>new-exp</i>—set the MPLS EXP bits, and send the packet. The range is 0 to 7. set-prec-transmit <i>new-prec</i>—set the IP precedence value to a new value, and send the packet. The range is 0 to 7. transmit—send the packet without altering it. If you set the CIR equal to the PIR, a traffic rate that is less than the CIR or that meets the CIR is in the conform range. Traffic that exceeds the CIR rate is in the violate range. If you set the PIR greater than the CIR, a traffic rate less than the CIR is in the conform range. A traffic rate that exceeds the CIR but is less than or equal to the PIR is in the exceed range. A traffic rate that exceeds the PIR is in the violate range. Setting the burst sizes too low can result in less traffic than expected, and setting them too high can result in more traffic than expected.
Step 5	exit	Return to policy-map configuration mode.

	Command	Purpose
Step 6	exit	Return to global configuration mode.
Step 7	interface <i>interface-id</i>	Specify an ES port to attach to the policy map, and enter interface configuration mode.
Step 8	service-policy output <i>policy-map-name</i>	Specify the egress policy-map name, and apply it to the ES port. Only one policy map per port is supported.
Step 9	end	Return to privileged EXEC mode.
Step 10	show policy-map <i>policy-map-name</i> or show policy-map interface <i>interface-id</i>	Verify your entries.
Step 11	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, use the **no class** *class-name* policy-map configuration command. To remove the two-rate policer, use the **no police cir** *cir* [**bc** *conform-burst*] **pir** *pir* [**be** *peak-burst*] [**conform-action** *action* [**exceed-action** *action* [**violate-action** *action*]]] policy-map class configuration command. To remove the policy map and interface association, use the **no service-policy output** *policy-map-name* interface configuration command.

This example shows how to configure a class-level, two-rate traffic policer to limit outbound traffic to an average committed rate of 500 kbps and a peak rate of 1 Mbps. Traffic marked as conforming to the average committed rate (500 kbps) is sent as is. Traffic marked as exceeding 500 kbps, but not exceeding 1 Mbps, is marked with IP precedence 2 and then sent. All traffic marked as exceeding 1 Mbps is dropped. The burst parameters are set to 10000 bytes.

```
Switch(config)# class-map c1
Switch(config-cmap)# match ip precedence 4
Switch(config-cmap)# exit
Switch(config)# policy-map policy1
Switch(config-pmap)# class c1
Switch(config-pmap-c)# police cir 500000 bc 10000 pir 1000000 be 10000 conform-action
transmit exceed-action set-prec-transmit 2 violate-action drop
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/1
Switch(config-if)# service-policy output policy1
```

This example shows how to configure a class-level, two-rate traffic policer that uses a CIR and a PIR based on a percentage of bandwidth. A CIR of 20 percent and a PIR of 40 percent are specified. The optional **bc** and **be** values (300 ms and 400 ms, respectively) are specified.

```
Switch(config)# class-map c1
Switch(config-cmap)# match ip precedence 4
Switch(config-cmap)# exit
Switch(config)# policy-map policy1
Switch(config-pmap)# class class1
Switch(config-pmap-c)# police cir percent 20 bc 300 ms pir percent 40 be 400 ms
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/1
Switch(config-if)# service-policy output policy1
```

This example shows how to configure a VLAN-level, two-rate traffic policer to limit outbound traffic to an average committed rate of 500 kbps and a peak rate of 1 Mbps. Traffic marked as conforming to the average committed rate (500 kbps) is sent as is. Traffic marked as exceeding 500 kbps, but not exceeding 1 Mbps, is marked with IP precedence 2 and then sent. All traffic marked as exceeding 1 Mbps is dropped. The burst parameters are set to 10000 bytes.

```
Switch(config)# class-map match-all vlan203
Switch(config-cmap)# match vlan 203
Switch(config-cmap)# match vlan inner 206
Switch(config-cmap)# exit
Switch(config)# policy-map vlan-policy
Switch(config-pmap)# class vlan203
Switch(config-pmap-c)# police cir 500000 bc 10000 pir 1000000 be 10000 conform-action
transmit exceed-action set-prec-transmit 2 violate-action drop
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/2
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport mode trunk
Switch(config-if)# service-policy output vlan-policy
```

This example shows how to create a hierarchical service-policy in which all levels are present. This configuration associates a class-level policy-map with a VLAN-level policy-map, associates the VLAN-level policy-map with a physical-level policy-map, and attaches the physical-level policy-map to a physical port.

Within a policy map, the class-default applies to all traffic that is not explicitly matched within the policy map but does match the parent policy. If no parent policy is configured, the parent policy represents the physical port. In a physical-level policy-map, class-default is the only class that can be configured.

```
Switch(config)# class-map my-class
Switch(config-cmap)# match ip precedence 1
Switch(config-cmap)# exit
Switch(config)# class-map my-logical-class
Switch(config-cmap)# match vlan 5
Switch(config-cmap)# exit
Switch(config)# policy-map my-class-policy
Switch(config-pmap)# class my-class
Switch(config-pmap-c)# set ip precedence 2
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# policy-map my-logical-policy
Switch(config-pmap)# class my-logical-class
Switch(config-pmap-c)# service-policy my-class-policy
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# policy-map my-physical-policy
Switch(config-pmap)# class class-default
Switch(config-pmap-c)# police cir 500000 bc 10000 pir 1000000 be 10000 conform-action
transmit exceed-action set-prec-transmit 2 violate-action drop
Switch(config-pmap-c)# service-policy my-logical-policy
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/1
Switch(config)# service-policy output my-physical-policy
```

Configuring Class-Based Packet Marking in an Egress Traffic Policy

You can perform class-based packet marking in a traffic policy by configuring the **set {cos new-cos | ip {dscp new-dscp | precedence new-precedence} | mpls experimental exp-number}** policy-map class configuration command in an egress policy-map attached to an ES port.

Before beginning this procedure, make sure that you have created the class map to isolate traffic. For more information, see the [“Classifying Egress Traffic by Using Class Maps”](#) section on page 26-78.

Beginning in privileged EXEC mode, follow these steps to configure class-level, class-based packet marking in a service policy for outbound traffic. This procedure is optional. The examples that follow the procedure show how to configure class-level and VLAN-level class-based packet marking.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	policy-map <i>policy-map-name</i>	Create a policy map by entering the policy-map name, and enter policy-map configuration mode. By default, no policy maps are defined.
Step 3	class <i>class-name</i>	Specify the name of the class whose traffic policy you want to create or change, and enter policy-map class configuration mode. By default, no traffic classes are defined.
Step 4	set {cos new-cos ip {dscp new-dscp precedence new-precedence} mpls experimental exp-number}	Designate the value to set in the traffic class if the packets match the criteria. - For cos new-cos, enter the new CoS value assigned to the classified traffic. The range is 0 to 7. - For ip dscp new-dscp, enter the new DSCP value assigned to the classified traffic. The range is 0 to 63. The specified value sets the ToS byte in the packet header. - For ip precedence new-precedence, enter the new IP-precedence value assigned to the classified traffic. The range is 0 to 7. The specified value sets the precedence bit in the IP header. - For mpls experimental exp-number, enter the new MPLS EXP value assigned to the classified traffic. The range is 0 to 7. The specified value sets the MPLS EXP 3-bit field in the packet header.
Step 5	exit	Return to policy-map configuration mode.
Step 6	exit	Return to global configuration mode.
Step 7	interface <i>interface-id</i>	Specify an ES port to attach to the policy map, and enter interface configuration mode.
Step 8	service-policy output <i>policy-map-name</i>	Specify the egress policy-map name, and apply it to the ES port. Only one policy map per port is supported.
Step 9	end	Return to privileged EXEC mode.

	Command	Purpose
Step 10	show policy-map [<i>policy-map-name</i> [class <i>class-map-name</i>]] or show policy-map interface <i>interface-id</i>	Verify your entries.
Step 11	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, use the **no class** *class-name* policy-map configuration command. To remove an assigned value, use the **no set** {**cos** *new-cos* | **ip** {**dscp** *new-dscp* | **precedence** *new-precedence*} | **mpls experimental** *exp-number*} policy-map class configuration command. To remove the policy map and interface association, use the **no service-policy output** *policy-map-name* interface configuration command.

This example shows how to create a class-level policy-map called *out_pmap*. When it is attached to the ES port, it matches packets with MPLS EXP field 2 and resets this field to 3.

```
Switch(config)# class-map mpls_2
Switch(config-cmap)# match mpls experimental 2
Switch(config-cmap)# exit
Switch(config)# policy-map out-pmap
Switch(config-pmap)# class mpls_2
Switch(config-pmap-c)# set mpls experimental 3
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/1
Switch(config-if)# service-policy output out-pmap
```

This example shows how to create a VLAN-level policy-map called *log-policy*. It matches packets with VLAN 203 and associates a class-level child policy called *cls-policy*. The child policy matches packets with MPLS EXP 2 and resets them to 5.

```
Switch(config)# class-map cls-class
Switch(config-cmap)# match mpls experimental 2
Switch(config-cmap)# exit
Switch(config)# class-map log-class
Switch(config-cmap)# match vlan 203
Switch(config-cmap)# exit
Switch(config)# policy-map cls-policy
Switch(config-pmap)# class cls-class
Switch(config-pmap-c)# set mpls experimental 5
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# policy-map log-policy
Switch(config-pmap)# class log-class
Switch(config-pmap-c)# service-policy cls-policy
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/2
Switch(config-if)# switchport trunk encapsulation isl
Switch(config-if)# switchport mode trunk
Switch(config-if)# service-policy output log-policy
```

Configuring CBWFQ and Tail Drop

CBWFQ creates a queue for every class for which a class map is defined. Packets satisfying the match criteria for a class accumulate in the queue reserved for the class until they are sent, which occurs when the queue is serviced by the fair queue process. When the maximum packet threshold that you defined for the class is reached, any more packets destined for the class queue are dropped according to the tail drop or WRED mechanism.

You configure tail drop by using the **queue-limit** policy-map class configuration command in an egress policy-map attached to an ES port. This command configures the maximum threshold for tail drop. Packets are queued until the maximum threshold is exceeded, and then all the packets are dropped.

Before beginning this procedure, make sure that you have reviewed the configuration guidelines and have created the class map to isolate traffic. For more information, see the [“Hierarchical QoS Configuration Guidelines” section on page 26-76](#) and the [“Classifying Egress Traffic by Using Class Maps” section on page 26-78](#). For information on how to configure WRED, see the [“Configuring CBWFQ and DSCP-Based WRED” section on page 26-89](#).

Beginning in privileged EXEC mode, follow these steps to configure class-level CBWFQ and tail drop in a service policy for outbound traffic. This procedure is optional. The examples that follow the procedure show how to configure class-level and VLAN-level CBWFQ and tail drop.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	policy-map <i>policy-map-name</i>	Create a policy map by entering the policy-map name, and enter policy-map configuration mode. By default, no policy maps are defined.
Step 3	class <i>class-name</i>	Specify the name of the class whose traffic policy you want to create or change, and enter policy-map class configuration mode. By default, no traffic classes are defined.

	Command	Purpose
Step 4	bandwidth { <i>bandwidth-kbps</i> percent <i>percent</i> }	Specify the minimum bandwidth provided to a class belonging to the egress policy-map when there is traffic congestion in the switch. If the switch is not congested, the class receives more bandwidth than you specify with the bandwidth command. CBWFQ derives the weight for packets belonging to the class from the bandwidth allocated to the class. CBWFQ then uses the weight to ensure that the queue for the class is serviced fairly. By default, no bandwidth is specified. You can specify the bandwidth in kbps or as a percentage: - For <i>bandwidth-kbps</i>, specify the bandwidth amount in kbps assigned to the class. The range is 200 to 2000000. Allocate the bandwidth in 100-kbps increments; otherwise, the software rounds down the bandwidth to the nearest 100-kbps increment. - For percent <i>percent</i>, specify the percentage of available bandwidth assigned to the class. The range is 1 to 100. The sum of the class bandwidth percentages within a single policy map cannot exceed 99 percent. Percentage calculations are based on the bandwidth available at the parent class (or the physical level if it is the parent). Specify all the class bandwidths in either kbps or in percentages, but not a mix of both. The amount of bandwidth configured should be large enough to accommodate Layer 2 overhead.
Step 5	queue-limit <i>limit</i>	Configure the maximum threshold for tail drop. For <i>limit</i>, the range is 1 to 32768 packets. The default is 128 packets.
Step 6	exit	Return to policy-map configuration mode.
Step 7	exit	Return to global configuration mode.
Step 8	interface <i>interface-id</i>	Specify an ES port to attach to the policy map, and enter interface configuration mode.
Step 9	service-policy output <i>policy-map-name</i>	Specify the egress policy-map name, and apply it to the ES port. Only one policy map per port is supported.
Step 10	end	Return to privileged EXEC mode.
Step 11	show policy-map [<i>policy-map-name</i> [class <i>class-map-name</i>]] or show policy-map interface <i>interface-id</i>	Verify your entries.
Step 12	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, use the **no class** *class-name* policy-map configuration command. To return to the default bandwidth, use the **no bandwidth** policy-map class configuration command. To return to the default maximum threshold, use the **no queue-limit** policy-map class configuration command.

This example shows how to create a class-level policy-map called *policy11* for three classes called *prec1*, *prec2*, and *prec3*. In the policy for these classes, 30 percent of the available bandwidth is assigned to the queue for the first class, 20 percent is assigned to the queue for the second class, and 10 percent is assigned to the queue for the third class. Tail drop is enabled on each class queue, packets are queued until the maximum threshold of 2000 packets is exceeded, and then all the packets are dropped.

```
Switch(config)# class-map prec1
Switch(config-cmap)# match ip precedence 1
Switch(config-cmap)# exit
Switch(config)# class-map prec2
Switch(config-cmap)# match ip precedence 2
Switch(config-cmap)# exit
Switch(config)# class-map prec3
Switch(config-cmap)# match ip precedence 3
Switch(config-cmap)# exit
Switch(config)# policy-map policy11
Switch(config-pmap)# class prec1
Switch(config-pmap-c)# bandwidth percent 30
Switch(config-pmap-c)# queue-limit 2000
Switch(config-pmap-c)# exit
Switch(config-pmap)# class prec2
Switch(config-pmap-c)# bandwidth percent 20
Switch(config-pmap-c)# queue-limit 2000
Switch(config-pmap-c)# exit
Switch(config-pmap)# class prec3
Switch(config-pmap-c)# bandwidth percent 10
Switch(config-pmap-c)# queue-limit 2000
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/1
Switch(config-if)# service-policy output policy11
```

This example shows how to create a VLAN-level policy-map called *vlan-policy* for two classes called *vlan203* and *vlan202*. In the policy for these classes, the minimum bandwidth is set to 2000 kbps.

```
Switch(config)# class-map match-all vlan203
Switch(config-cmap)# match vlan 203
Switch(config-cmap)# match vlan inner 206
Switch(config-cmap)# exit
Switch(config)# class-map match-all vlan202
Switch(config-cmap)# match vlan 202
Switch(config-cmap)# match vlan inner 204
Switch(config-cmap)# exit
Switch(config)# policy-map vlan-policy
Switch(config-pmap)# class vlan203
Switch(config-pmap-c)# bandwidth 2000
Switch(config-pmap-c)# exit
Switch(config-pmap)# class vlan202
Switch(config-pmap-c)# bandwidth 2000
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/2
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport mode trunk
Switch(config-if)# service-policy output vlan-policy
```

This example shows how to create a VLAN-level policy-map called *log-policy*. It matches packets with VLAN 203 and associates a class-level child policy called *cls-policy*. The child policy matches packets with a CoS value of 2, sets 50 percent as the available bandwidth for this class, and sets 2000 packets as the maximum threshold for tail drop. Packets are queued until the maximum threshold is exceeded, and then all the packets are dropped. Note that when you configure the **bandwidth** command in a class policy, you also must configure the **bandwidth** or **shape** policy-map class configuration command in the parent VLAN-level policy.

```
Switch(config)# class-map cls-class
Switch(config-cmap)# match cos 2
Switch(config-cmap)# exit
Switch(config)# class-map log-class
Switch(config-cmap)# match vlan 203
Switch(config-cmap)# exit
Switch(config)# policy-map cls-policy
Switch(config-pmap)# class cls-class
Switch(config-pmap-c)# bandwidth percent 50
Switch(config-pmap-c)# queue-limit 2000
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# policy-map log-policy
Switch(config-pmap)# class log-class
Switch(config-pmap-c)# bandwidth percent 50
Switch(config-pmap-c)# service-policy cls-policy
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/2
Switch(config-if)# switchport trunk encapsulation isl
Switch(config-if)# switchport mode trunk
Switch(config-if)# service-policy output log-policy
```

Configuring CBWFQ and DSCP-Based WRED

CBWFQ creates a queue for every class for which a class map is defined. Packets satisfying the match criteria for a class accumulate in the queue reserved for the class until they are sent, which occurs when the queue is serviced by the fair queue process. When the maximum packet threshold that you defined for the class is reached, any more packets destined for the class queue are dropped according to the tail drop or the WRED mechanism.

WRED reduces the chances of tail drop by selectively dropping packets when the port begins to show signs of congestion. By dropping some packets early rather than waiting until the queue is full, WRED avoids dropping large numbers of packets at once.

When a packet arrives and WRED is enabled, these events occur:

- The average queue size is calculate based on the previous average and the current size of the queue. The average queue-size calculation is affected by the exponential weight constant setting in the **random-detect exponential-weight-constant** policy-map class configuration command.
- If the average queue size is less than the minimum queue threshold, the arriving packet is queued. The minimum queue threshold is configured through the *min-threshold* option in the **random-detect dscp** policy-map class configuration command.

- If the average queue size is between the minimum queue threshold and the maximum queue threshold, the packet is either dropped or queued, depending on the packet-drop probability. The packet-drop probability is based on the minimum threshold, the maximum threshold, and the mark-probability denominator. The maximum queue threshold is configured through the *max-threshold* option, and the mark-probability denominator is configured through the *mark-prob-denominator* option in the **random-detect dscp** policy-map class configuration command.
- If the average queue size is greater than the maximum queue threshold, the packet is automatically dropped.

You enable DSCP-based WRED by using the **random-detect dscp-based** policy-map class configuration command in an egress policy-map attached to an ES port. This command allows for preferential drop treatment among packets with different DSCP values. The WRED algorithm discards or marks packets destined for a queue when that queue is congested. It discards packets fairly and before the queue is full. If you want to enable IP precedence-based WRED instead of DSCP-based WRED, see the [“Configuring CBWFQ and IP Precedence-Based WRED”](#) section on page 26-93.

Before beginning this procedure, make sure that you have reviewed the configuration guidelines and have created the class map to isolate traffic. For more information, see the [“Hierarchical QoS Configuration Guidelines”](#) section on page 26-76 and the [“Classifying Egress Traffic by Using Class Maps”](#) section on page 26-78. For information on how to configure tail drop, see the [“Configuring CBWFQ and Tail Drop”](#) section on page 26-86.

Beginning in privileged EXEC mode, follow these steps to configure class-level CBWFQ and DSCP-based WRED in a service policy for outbound traffic. This procedure is optional. The examples that follow the procedure show how to configure class-level and VLAN-level CBWFQ and DSCP-based WRED.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	policy-map <i>policy-map-name</i>	Create a policy map by entering the policy-map name, and enter policy-map configuration mode. By default, no policy maps are defined.
Step 3	class <i>class-name</i>	Specify the name of the class whose traffic policy you want to create or change, and enter policy-map class configuration mode. By default, no traffic classes are defined.

	Command	Purpose
Step 4	bandwidth { <i>bandwidth-kbps</i> percent <i>percent</i> }	Specify the minimum bandwidth provided to a class belonging to the egress policy-map when there is traffic congestion in the switch. If the switch is not congested, the class receives more bandwidth than you specify with the bandwidth command. CBWFQ derives the weight for packets belonging to the class from the bandwidth allocated to the class. CBWFQ then uses the weight to ensure that the queue for the class is serviced fairly. By default, no bandwidth is specified. You can specify the bandwidth in kbps or as a percentage: - For <i>bandwidth-kbps</i>, specify the bandwidth amount in kbps assigned to the class. The range is 200 to 2000000. Allocate the bandwidth in 100-kbps increments; otherwise, the software rounds down the bandwidth to the nearest 100-kbps increment. - For percent <i>percent</i>, specify the percentage of available bandwidth assigned to the class. The range is 1 to 100. The sum of the class bandwidth percentages within a single policy map cannot exceed 99 percent. Percentage calculations are based on the bandwidth available at the parent class (or the physical level if it is the parent). Specify all the class bandwidths in either kbps or in percentages, but not a mix of both. The amount of bandwidth configured should be large enough to accommodate Layer 2 overhead.
Step 5	random-detect dscp-based	Enable DSCP-based WRED as a drop policy. By default, WRED is disabled.
Step 6	random-detect dscp <i>dscp min-threshold max-threshold mark-prob-denominator</i>	Specify packet threshold and mark-probability values for a specific DSCP value: - For <i>dscp</i>, enter a DSCP value. The range is 0 to 63. You also can enter a mnemonic name for a commonly used value. - For <i>min-threshold</i>, enter the minimum threshold in packets. The range is 1 to 32768. When the average queue size reaches the minimum threshold, WRED randomly drops some packets with the specified DSCP value. - For <i>max-threshold</i>, enter the maximum threshold in packets. The range is 1 to 32768. The default is 128. When the average queue size exceeds the maximum threshold, WRED drops all packets with the specified DSCP value. - For <i>mark-prob-denominator</i>, enter the denominator for the fraction of packets dropped when the average queue size is at the maximum threshold. The range is 1 to 65535. The default is 10. For example, if the denominator is 512, one out of every 512 packets is dropped when the queue is at the maximum threshold. For a list of the default settings for a specified DSCP value, see the command reference for this release.

	Command	Purpose
Step 7	random-detect exponential-weighting-constant <i>weight</i>	(Optional) Configure the exponential weight factor for the average queue-size calculation for WRED. The range is 1 to 16. The default is 9.
Step 8	exit	Return to policy-map configuration mode.
Step 9	exit	Return to global configuration mode.
Step 10	interface <i>interface-id</i>	Specify an ES port to attach to the policy map, and enter interface configuration mode.
Step 11	service-policy output <i>policy-map-name</i>	Specify the egress policy-map name, and apply it to the ES port. Only one policy map per port is supported.
Step 12	end	Return to privileged EXEC mode.
Step 13	show policy-map [<i>policy-map-name</i> [class <i>class-map-name</i>]] or show policy-map interface <i>interface-id</i>	Verify your entries.
Step 14	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, use the **no class** *class-name* policy-map configuration command. To return to the default bandwidth, use the **no bandwidth** policy-map class configuration command. To disable DSCP-based WRED, use the **no random-detect dscp-based** policy-map class configuration command. To return to the default WRED settings, use the **no random-detect dscp** *dscp* policy-map class configuration command.

This example shows how to configure a class-level policy called *policy10*. Class *c1* has these characteristics: a minimum of 2000 kbps of bandwidth is expected to be delivered to this class in the event of congestion, and a weight factor of 10 is used to calculate the average queue size. To avoid congestion, DSCP-based WRED packet drop is used instead of tail drop. The minimum threshold for DSCP value 8 is 24, the maximum threshold is 40, and the mark-probability denominator is 512.

```
Switch(config)# class-map c1
Switch(config-cmap)# match ip dscp 8
Switch(config-cmap)# exit
Switch(config)# policy-map policy10
Switch(config-pmap)# class c1
Switch(config-pmap-c)# bandwidth 2000
Switch(config-pmap-c)# random-detect dscp-based
Switch(config-pmap-c)# random-detect exponential-weighting-constant 10
Switch(config-pmap-c)# random-detect dscp 8 24 40 512
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/1
Switch(config-if)# switchport trunk encapsulation isl
Switch(config-if)# switchport mode trunk
Switch(config-if)# service-policy output policy10
```


This example shows how to configure a VLAN-level policy called *parent*. It matches packets with VLAN 101 and associates a class-level child policy called *policy1*. The child policy matches two DSCP values in two classes. Thirty percent of the available bandwidth is assigned to the gold class, 20 percent is assigned to the silver class, and 20 percent is assigned to *vlan101*. DSCP-based WRED is used as the drop policy. For the *af11* value in the gold class, WRED randomly drops packets with this DSCP when the minimum threshold reaches 30. When the average queue size exceeds the maximum threshold of 40, WRED drops all packets with DSCP *af11*. The mark-probability denominator is set to 10, which means that one out of every 10 packets is dropped when the average queue is at the maximum threshold. The configuration has similar settings for *af12* in the gold class and for *af21* and *af22* in the silver class. Note that when you configure the **bandwidth** command in a class policy, you also must configure the **bandwidth** or **shape** policy-map class configuration command in the parent VLAN-level policy.

```
Switch(config)# class-map gold
Switch(config-cmap)# match ip dscp af11 af12
Switch(config-cmap)# exit
Switch(config)# class-map silver
Switch(config-cmap)# match ip dscp af21 af22
Switch(config-cmap)# exit
Switch(config)# class-map vlan101
Switch(config-cmap)# match vlan 101
Switch(config-cmap)# exit
Switch(config)# policy-map policy1
Switch(config-pmap)# class gold
Switch(config-pmap-c)# bandwidth percent 30
Switch(config-pmap-c)# random-detect dscp-based
Switch(config-pmap-c)# random-detect dscp af11 30 40 10
Switch(config-pmap-c)# random-detect dscp af12 25 40 10
Switch(config-pmap-c)# exit
Switch(config-pmap)# class silver
Switch(config-pmap-c)# bandwidth percent 20
Switch(config-pmap-c)# random-detect dscp-based
Switch(config-pmap-c)# random-detect dscp af21 28 35 10
Switch(config-pmap-c)# random-detect dscp af22 26 35 10
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# policy-map parent
Switch(config-pmap)# class vlan101
Switch(config-pmap-c)# bandwidth percent 30
Switch(config-pmap-c)# service-policy policy1
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/1
Switch(config-if)# service-policy output parent
```

Configuring CBWFQ and IP Precedence-Based WRED

CBWFQ creates a queue for every class for which a class map is defined. Packets satisfying the match criteria for a class accumulate in the queue reserved for the class until they are sent, which occurs when the queue is serviced by the fair queue process. When the maximum packet threshold you defined for the class is reached, any more packets destined for the class queue are dropped according to the tail drop or the WRED mechanism.

WRED reduces the chances of tail drop by selectively dropping packets when the port begins to show signs of congestion. By dropping some packets early rather than waiting until the queue is full, WRED avoids dropping large numbers of packets at once.

You enable IP precedence-based WRED by using the **random-detect precedence-based** policy-map class configuration command in an egress policy-map attached to an ES port. This command allows for preferential drop treatment among packets with different IP precedence values. The WRED algorithm

discards or marks packets destined for a queue when that queue is congested. It discards packets fairly and before the queue is full. Packets with high IP-precedence values are preferred over packets with low IP-precedence values. If you want to enable DSCP-based WRED instead of IP precedence-based WRED, see the “[Configuring CBWFQ and DSCP-Based WRED](#)” section on page 26-89

Before beginning this procedure, make sure that you have reviewed the configuration guidelines and have created the class map to isolate traffic. For more information, see the “[Hierarchical QoS Configuration Guidelines](#)” section on page 26-76 and the “[Classifying Egress Traffic by Using Class Maps](#)” section on page 26-78. For information on how to configure tail drop, see the “[Configuring CBWFQ and Tail Drop](#)” section on page 26-86.

Beginning in privileged EXEC mode, follow these steps to configure class-level CBWFQ and IP precedence-based WRED in a service policy for outbound traffic. This procedure is optional. The examples that follow the procedure show how to configure class-level and VLAN-level CBWFQ and IP precedence-based WRED.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	policy-map <i>policy-map-name</i>	Create a policy map by entering the policy-map name, and enter policy-map configuration mode. By default, no policy maps are defined.
Step 3	class <i>class-name</i>	Specify the name of the class whose traffic policy you want to create or change, and enter policy-map class configuration mode. By default, no traffic classes are defined.
Step 4	bandwidth { <i>bandwidth-kbps</i> percent <i>percent</i> }	Specify the minimum bandwidth provided to a class belonging to the egress policy-map when there is traffic congestion in the switch. If the switch is not congested, the class receives more bandwidth than you specify with the bandwidth command. CBWFQ derives the weight for packets belonging to the class from the bandwidth allocated to the class. CBWFQ then uses the weight to ensure that the queue for the class is serviced fairly. By default, no bandwidth is specified. You can specify the bandwidth in kbps or as a percentage: - For <i>bandwidth-kbps</i>, specify the bandwidth amount in kbps assigned to the class. The range is 200 to 2000000. Allocate the bandwidth in 100-kbps increments; otherwise, the software rounds down the bandwidth to the nearest 100-kbps increment. - For percent <i>percent</i>, specify the percentage of available bandwidth assigned to the class. The range is 1 to 100. The sum of the class bandwidth percentages within a single policy map cannot exceed 99 percent. Percentage calculations are based on the bandwidth available at the parent class (or the physical level if it is the parent). Specify all the class bandwidths in either kbps or in percentages, but not a mix of both. The amount of bandwidth configured should be large enough to accommodate Layer 2 overhead.
Step 5	random-detect precedence-based	Enable IP precedence-based WRED as a drop policy. By default, WRED is disabled.

	Command	Purpose
Step 6	random-detect precedence <i>ip-precedence</i> <i>min-threshold max-threshold</i> <i>mark-prob-denominator</i>	Specify packet threshold and mark-probability values for a specific IP precedence value: - For <i>ip-precedence</i>, specify an IP precedence value. The range is 0 to 7. - For <i>min-threshold</i>, specify the minimum threshold in packets. The range is 1 to 32768. When the average queue size reaches the minimum threshold, WRED randomly drops some packets with the specified precedence value. - For <i>max-threshold</i>, specify the maximum threshold in packets. The range is 1 to 32768. The default is 128. When the average queue size exceeds the maximum threshold, WRED drops all packets with the specified precedence value. - For <i>mark-prob-denominator</i>, specify the denominator for the fraction of packets dropped when the average queue size is at the maximum threshold. The range is 1 to 65535. The default is 10. For example, if the denominator is 512, one out of every 512 packets is dropped when the queue is at the maximum threshold. For a list of the default settings for a specified IP precedence value, see the command reference for this release.
Step 7	random-detect exponential-weighting-constant <i>weight</i>	(Optional) Configure the exponential weight factor for the average queue-size calculation for WRED. The range is 1 to 16. The default is 9.
Step 8	exit	Return to policy-map configuration mode.
Step 9	exit	Return to global configuration mode.
Step 10	interface <i>interface-id</i>	Specify an ES port to attach to the policy map, and enter interface configuration mode.
Step 11	service-policy output <i>policy-map-name</i>	Specify the egress policy-map name, and apply it to the ES port. Only one policy map per port is supported.
Step 12	end	Return to privileged EXEC mode.
Step 13	show policy-map [<i>policy-map-name</i> [class <i>class-map-name</i>]] or show policy-map interface <i>interface-id</i>	Verify your entries.
Step 14	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, use the **no class** *class-name* policy-map configuration command. To return to the default bandwidth, use the **no bandwidth** policy-map class configuration command. To disable IP precedence-based WRED, use the **no random-detect precedence-based** policy-map class configuration command. To return to the default WRED settings, use the **no random-detect precedence** *ip-precedence* policy-map class configuration command.

This example shows how to configure a class-level policy called *policy10*. Class *c1* has these characteristics: a minimum of 2000 kbps of bandwidth is expected to be delivered to this class in the event of congestion, and a weight factor of 10 is used to calculate the average queue size. To avoid congestion, WRED packet drop is used instead of tail drop. IP precedence is reset for levels 0 to 2. The minimum threshold for IP precedence value 0 is 32, the maximum threshold is 256, and the mark-probability denominator is 100. IP precedence values 1 and 2 have similar thresholds and probability denominators.

```
Switch(config)# class-map c1
Switch(config-cmap)# match ip precedence 0 1 2
Switch(config-cmap)# exit
Switch(config)# policy-map policy10
Switch(config-pmap)# class c1
Switch(config-pmap-c)# bandwidth 2000
Switch(config-pmap-c)# random-detect precedence-based
Switch(config-pmap-c)# random-detect exponential-weighting-constant 10
Switch(config-pmap-c)# random-detect precedence 0 32 256 100
Switch(config-pmap-c)# random-detect precedence 1 64 256 100
Switch(config-pmap-c)# random-detect precedence 2 96 256 100
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/1
Switch(config-if)# switchport trunk encapsulation isl
Switch(config-if)# switchport mode trunk
Switch(config-if)# service-policy output policy10
```

This example shows how to configure a VLAN-level policy called *parent*. It matches packets with VLAN 101 and associates a class-level child policy called *policy1*. The child policy matches two IP precedence values in two classes. Thirty percentage of the available bandwidth is assigned to the gold class, 20 percent is assigned to the silver class, and 30 percent is assigned to *vlan101*. IP precedence-based WRED is used as the drop policy. For the IP precedence 0 in the gold class, WRED randomly drops packets with this IP precedence when the minimum threshold reaches 30. When the average queue size exceeds the maximum threshold of 40, WRED drops all packets with IP precedence 0. The mark-probability denominator is set to 10, which means that one out of every 10 packets is dropped when the average queue is at the maximum threshold. The configuration has similar settings for IP precedence 3 in the silver class. Note that when you configure the **bandwidth** command in a class policy, you also must configure the **bandwidth** or **shape** policy-map class configuration command in the parent VLAN-level policy.

```
Switch(config)# class-map gold
Switch(config-cmap)# match ip precedence 0
Switch(config-cmap)# exit
Switch(config)# class-map silver
Switch(config-cmap)# match ip precedence 3
Switch(config-cmap)# exit
Switch(config)# class-map vlan101
Switch(config-cmap)# match vlan 101
Switch(config-cmap)# exit
Switch(config)# policy-map policy1
Switch(config-pmap)# class gold
Switch(config-pmap-c)# bandwidth percent 30
Switch(config-pmap-c)# random-detect precedence-based
Switch(config-pmap-c)# random-detect precedence 0 30 40 10
Switch(config-pmap-c)# exit
Switch(config-pmap)# class silver
Switch(config-pmap-c)# bandwidth percent 20
Switch(config-pmap-c)# random-detect precedence-based
Switch(config-pmap-c)# random-detect precedence 3 28 35 10
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
```

```
Switch(config)# policy-map parent
Switch(config-pmap)# class vlan101
Switch(config-pmap-c)# bandwidth percent 30
Switch(config-pmap-c)# service-policy policy1
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/1
Switch(config-if)# service-policy output parent
```

Enabling LLQ

LLQ provides strict-priority queueing for a traffic class. It enables delay-sensitive data, such as voice, to be sent before packets in other queues are sent. The priority queue is serviced first until it is empty. Only one traffic stream can be destined for the priority queue per class-level policy. The priority queue restricts all traffic streams in the same hierarchy, and you should use care when configuring this feature. You enable the priority queue for a traffic class by using the **priority** policy-map class configuration command in an egress policy-map attached to an ES port.

Before beginning this procedure, make sure that you have reviewed the configuration guidelines and have created the class map to isolate traffic. For more information, see the [“Hierarchical QoS Configuration Guidelines”](#) section on page 26-76 and the [“Classifying Egress Traffic by Using Class Maps”](#) section on page 26-78.

Beginning in privileged EXEC mode, follow these steps to enable class-level priority queueing in a service policy for outbound traffic. This procedure is optional. The examples that follow the procedure show how to enable class-level and VLAN-level priority queueing.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	policy-map <i>policy-map-name</i>	Create a policy map by entering the policy-map name, and enter policy-map configuration mode. By default, no policy maps are defined.
Step 3	class <i>class-name</i>	Specify the name of the class whose traffic policy you want to create or change, and enter policy-map class configuration mode. By default, no traffic classes are defined.
Step 4	priority	Enable the strict priority queue, and give priority to a class of traffic. By default, strict priority queueing is disabled.
Step 5	exit	Return to policy-map configuration mode.
Step 6	exit	Return to global configuration mode.
Step 7	interface <i>interface-id</i>	Specify an ES port to attach to the policy map, and enter interface configuration mode.
Step 8	service-policy output <i>policy-map-name</i>	Specify the egress policy-map name, and apply it to the ES port. Only one policy map per port is supported.
Step 9	end	Return to privileged EXEC mode.

	Command	Purpose
Step 10	show policy-map [<i>policy-map-name</i> [class <i>class-map-name</i>]] or show policy-map interface <i>interface-id</i>	Verify your entries.
Step 11	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, use the **no class** *class-name* policy-map configuration command. To disable the priority queue, use the **no priority** policy-map class configuration command.

This example shows how to configure a class-level policy called *policy1* that has two classes. Class 1 has 10 percent of the available bandwidth. Class 2 is configured as the priority queue, which is serviced first until it is empty.

```
Switch(config)# class-map class1
Switch(config-cmap)# match mpls experimental 2 3 4
Switch(config-cmap)# exit
Switch(config)# class-map class2
Switch(config-cmap)# match mpls experimental 7
Switch(config-cmap)# exit
Switch(config)# policy-map policy1
Switch(config-pmap)# class class1
Switch(config-pmap-c)# bandwidth percent 10
Switch(config-pmap-c)# exit
Switch(config-pmap)# class class2
Switch(config-pmap-c)# priority
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/2
Switch(config-if)# service-policy output policy1
```

This example shows how to configure a VLAN-level policy called *parent*. It matches packets with VLAN 101 and associates a class-level child policy called *policy1*. The gold class is configured as the priority queue, whereas the silver class and vlan101 class each have 20 percent of the available bandwidth assigned to them. DSCP-based WRED is used as the drop policy. For the af21 value in the silver class, WRED randomly drops packets with this DSCP when the minimum threshold reaches 28. When the average queue size exceeds the maximum threshold of 35, WRED drops all packets with DSCP af21. The mark-probability denominator is set to 10, which means that one out of every 10 packets is dropped when the average queue is at the maximum threshold. The configuration has similar settings for af22 in the silver class.

```
Switch(config)# class-map gold
Switch(config-cmap)# match ip dscp af11 af12
Switch(config-cmap)# exit
Switch(config)# class-map silver
Switch(config-cmap)# match ip dscp af21 af22
Switch(config-cmap)# exit
Switch(config)# class-map vlan101
Switch(config-cmap)# match vlan 101
Switch(config-cmap)# exit
Switch(config)# policy-map policy1
Switch(config-pmap)# class gold
Switch(config-pmap-c)# priority
Switch(config-pmap-c)# exit
Switch(config-pmap)# class silver
Switch(config-pmap-c)# bandwidth percent 20
Switch(config-pmap-c)# random-detect dscp-based
```

```

Switch(config-pmap-c)# random-detect dscp af21 28 35 10
Switch(config-pmap-c)# random-detect dscp af22 26 35 10
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# policy-map parent
Switch(config-pmap)# class vlan101
Switch(config-pmap-c)# bandwidth percent 20
Switch(config-pmap-c)# service-policy policy1
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/1
Switch(config-if)# service-policy output parent

```

Configuring Shaping

Shaping provides a process for delaying out-of-profile packets in queues so that they conform to a specified profile. Shaping is distinct from policing. Policing drops packets that exceed a configured threshold, but shaping buffers packets so that traffic remains within a threshold. Shaping offers greater smoothness in handling traffic than policing. You can configure average-rate traffic shaping on a traffic class within a policy map at the class level, at the VLAN level, and at the physical level by using the **shape** policy-map class configuration command. At the physical level of the hierarchy, you can shape only the class-default class in an egress policy attached to an ES port.

Before beginning this procedure, make sure that you have reviewed the configuration guidelines and have created the class map to isolate traffic. For more information, see the [“Hierarchical QoS Configuration Guidelines” section on page 26-76](#) and the [“Classifying Egress Traffic by Using Class Maps” section on page 26-78](#).

Beginning in privileged EXEC mode, follow these step to configure class-level shaping in a service policy for outbound traffic. This procedure is optional. The examples that follow the procedure show how to configure class-level, VLAN-level, and physical-level shaping.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	policy-map <i>policy-map-name</i>	Create a policy map by entering the policy-map name, and enter policy-map configuration mode. By default, no policy maps are defined.
Step 3	class <i>class-name</i>	Specify the name of the class whose traffic policy you want to create or change, and enter policy-map class configuration mode. By default, no traffic classes are defined.
Step 4	shape average <i>cir-bps</i>	Enable average-rate traffic shaping. Specify the committed information rate, the bit rate that traffic is shaped to, in bps. This is the access bit rate that you have contracted for with your service provider or the service levels that you intend to maintain. The range is 64000 to 2000000000 bps. Allocate the shaped rate in 100-kbps increments; otherwise, the software rounds down the bandwidth to the nearest 100-kbps increment. By default, average-rate traffic shaping is disabled.
Step 5	exit	Return to policy-map configuration mode.
Step 6	exit	Return to global configuration mode.

	Command	Purpose
Step 7	interface <i>interface-id</i>	Specify an ES port to attach to the policy map, and enter interface configuration mode.
Step 8	service-policy output <i>policy-map-name</i>	Specify the egress policy-map name, and apply it to the ES port. Only one policy map per port is supported.
Step 9	end	Return to privileged EXEC mode.
Step 10	show policy-map [<i>policy-map-name</i> [class <i>class-map-name</i>]] or show policy-map interface <i>interface-id</i>	Verify your entries.
Step 11	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, use the **no class** *class-name* policy-map configuration command. To disable the average-rate traffic shaping, use the **no shape average** policy-map class configuration command.

This example shows how to configure class-level, average-rate shaping. It limits traffic class *class1* to a data transmission rate of 256 kbps.

```
Switch(config)# class-map class1
Switch(config-cmap)# match cos 0 1 2 3
Switch(config-cmap)# exit
Switch(config)# policy-map policy1
Switch(config-pmap)# class class1
Switch(config-pmap-c)# shape average 256000
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/1
Switch(config-if)# service-policy output policy1
```

This example shows how to configure VLAN-level, average-rate shaping. It limits each traffic class, *vlan101* and *vlan102*, to a data transmission rate of 400 Mbps.

```
Switch(config)# class-map match-all vlan101
Switch(config-cmap)# match vlan 101
Switch(config-cmap)# exit
Switch(config)# class-map match-all vlan102
Switch(config-cmap)# match vlan 102
Switch(config-cmap)# exit
Switch(config)# policy-map vlan-policy
Switch(config-pmap)# class vlan101
Switch(config-pmap-c)# shape average 400000000
Switch(config-pmap-c)# exit
Switch(config-pmap)# class vlan102
Switch(config-pmap-c)# shape average 400000000
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet1/1/1
Switch(config-if)# switchport trunk encapsulation isl
Switch(config-if)# switchport mode trunk
Switch(config-if)# service-policy output vlan-policy
```


This example shows how to shape the default class. This configuration associates a class-level policy-map with a VLAN-level policy-map and then associates the VLAN-level policy-map with a physical-level policy-map.

```
Switch(config)# class-map my-class
Switch(config-cmap)# match ip precedence 1
Switch(config-cmap)# exit
Switch(config)# class-map my-logical-class
Switch(config-cmap)# match vlan 5
Switch(config-cmap)# exit
Switch(config)# policy-map my-class-policy
Switch(config-pmap)# class my-class
Switch(config-pmap-c)# set ip precedence 2
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# policy-map my-logical-policy
Switch(config-pmap)# class my-logical-class
Switch(config-pmap-c)# shape average 400000000
Switch(config-pmap-c)# service-policy my-class-policy
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# policy-map my-physical-policy
Switch(config-pmap)# class class-default
Switch(config-pmap-c)# shape average 500000000
Switch(config-pmap-c)# service-policy my-logical-policy
```

Displaying Hierarchical QoS Information

To display hierarchical QoS information, use one or more of the privileged EXEC commands in [Table 26-16](#):

Table 26-16 Commands for Displaying Hierarchical 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 mls qos	Display global QoS configuration information.
show policy-map [<i>policy-map-name</i> [<i>class class-map-name</i>]]	Display QoS policy-maps, which define the traffic policy for a traffic class.
show policy-map interface <i>interface-id</i> output [<i>class class-name</i>]]	Display QoS policy-map information for the specified ES port, and display statistics for an individual class.



Configuring EtherChannels

This chapter describes how to configure EtherChannels on Layer 2 and Layer 3 ports on the Catalyst 3750 Metro 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

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 EtherChannels, page 27-1](#)
- [Configuring EtherChannels, page 27-8](#)
- [Displaying EtherChannel, PAgP, and LACP Status, page 27-19](#)

Understanding EtherChannels


These sections describe how EtherChannels work:

- [EtherChannel Overview, page 27-2](#)
- [Port-Channel Interfaces, page 27-3](#)
- [Port Aggregation Protocol, page 27-4](#)
- [Link Aggregation Control Protocol, page 27-5](#)
- [Load Balancing and Forwarding Methods, page 27-6](#)

EtherChannel Overview

An EtherChannel consists of individual Fast Ethernet or Gigabit Ethernet links bundled into a single logical link as shown in [Figure 27-1](#).

Figure 27-1 Typical EtherChannel Configuration



The EtherChannel provides full-duplex bandwidth up to 800 Mbps (Fast EtherChannel) or 2 Gbps (Gigabit EtherChannel) between your switch and another switch or host. You cannot bundle standard Gigabit Ethernet ports with the enhanced-services (ES) Gigabit Ethernet ports.

Each EtherChannel can consist of up to eight compatibly configured Ethernet ports. All ports in each EtherChannel must be configured as either Layer 2 or Layer 3 ports. For Catalyst 3750 Metro switches, the number of EtherChannels is limited to 12. For more information, see the [“EtherChannel Configuration Guidelines”](#) section on page 27-9. The EtherChannel Layer 3 ports are made up of routed ports. Routed ports are physical ports configured to be in Layer 3 mode by using the **no switchport** interface configuration command. For more information, see the [Chapter 9, “Configuring Interface Characteristics.”](#)

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.

Port-Channel Interfaces

When you create an EtherChannel, a port-channel logical interface is involved:

- With Layer 2 ports, use the **channel-group** interface configuration command to dynamically create the port-channel logical interface.


You also can use the **interface port-channel** *port-channel-number* global configuration command to manually create the port-channel logical interface, but then you must use the **channel-group** *channel-group-number* command to bind the logical interface to a physical port. The *channel-group-number* can be the same as the *port-channel-number*, or you can use a new number. If you use a new number, the **channel-group** command dynamically creates a new port channel.

- With Layer 3 ports, you should manually create the logical interface by using the **interface port-channel** global configuration command followed by the **no switchport** interface configuration command. Then you manually assign an interface to the EtherChannel by using the **channel-group** interface configuration command.

For both Layer 2 and Layer 3 ports, the **channel-group** command binds the physical port and the logical interface together as shown in Figure 27-2.

Each EtherChannel has a port-channel logical interface numbered from 1 to 12. This port-channel interface number corresponds to the one specified with the **channel-group** interface configuration command.

Figure 27-2 Relationship of Physical Ports, Logical Port Channels, and Channel Groups



After you configure an EtherChannel, configuration changes applied to the port-channel apply to all the physical ports assigned to the port-channel. Configuration changes applied to the physical port affect only the port where you apply the configuration. To change the parameters of all ports in an EtherChannel, apply configuration commands to the port-channel interface, for example, spanning-tree commands or commands to configure a Layer 2 EtherChannel as a trunk.

Port Aggregation Protocol

The Port Aggregation Protocol (PAgP) is a Cisco-proprietary protocol that can be run only on Cisco switches and on those switches licensed by vendors to support PAgP. PAgP facilitates the automatic creation of EtherChannels by exchanging PAgP packets between Ethernet ports.

By using PAgP, the switch learns the identity of partners capable of supporting PAgP and the capabilities of each port. It then dynamically groups similarly configured ports into a single logical link (channel or aggregate port). Similarly configured ports are grouped based on hardware, administrative, and port parameter constraints. For example, PAgP groups the ports 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 Modes

Table 27-1 shows the user-configurable EtherChannel PAgP modes for the **channel-group** interface configuration command.

Table 27-1 EtherChannel PAgP Modes

Mode	Description
auto	Places a port into a passive negotiating state, in which the port responds to PAgP packets it receives but does not start PAgP packet negotiation. This setting minimizes the transmission of PAgP packets.
desirable	Places a port into an active negotiating state, in which the port starts negotiations with other ports by sending PAgP packets.
on	Forces a port to channel without PAgP (or the Link Aggregation Control Protocol [LACP]). With the on mode, a usable EtherChannel exists only when a port group in the on mode is connected to another port group in the on mode.

Switch ports exchange PAgP packets only with partner ports configured in the **auto** or **desirable** modes. Ports configured in the **on** mode do not exchange PAgP packets.

Both the **auto** and **desirable** modes enable ports to negotiate with partner ports to form an EtherChannel based on criteria such as port speed and, for Layer 2 EtherChannels, trunking state and VLAN numbers.

Ports can form an EtherChannel when they are in different PAgP modes as long as the modes are compatible. For example:

- A port in the **desirable** mode can form an EtherChannel with another port that is in the **desirable** or **auto** mode.
- A port in the **auto** mode can form an EtherChannel with another port in the **desirable** mode.

A port in the **auto** mode cannot form an EtherChannel with another port that is also in the **auto** mode because neither port starts PAgP negotiation.

A port in the **on** mode that is added to a port channel is forced to have the same characteristics as the already existing **on** mode ports in the channel.



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.

If your switch is connected to a partner that is PAgP-capable, you can configure the switch port 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.

Use the silent mode 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 port to a channel group, and to use the port for transmission.

PAgP Interaction with Other Features

The Dynamic Trunking Protocol (DTP) and the Cisco Discovery Protocol (CDP) send and receive packets over the physical ports in the EtherChannel. Trunk ports send and receive PAgP protocol data units (PDUs) on the lowest numbered VLAN.

In Layer 2 EtherChannels, the first port in the channel that comes up provides its MAC address to the EtherChannel. If this port is removed from the bundle, one of the remaining ports in the bundle provides its MAC address to the EtherChannel.

PAgP sends and receives PAgP PDUs only from ports that are up and have PAgP enabled for the auto or desirable mode.

Link Aggregation Control Protocol

The LACP is defined in IEEE 802.3ad and enables Cisco switches to manage Ethernet channels between switches that conform to the 802.3ad protocol. LACP facilitates the automatic creation of EtherChannels by exchanging LACP packets between Ethernet ports.

By using LACP, the switch learns the identity of partners capable of supporting LACP and the capabilities of each port. It then dynamically groups similarly configured ports into a single logical link (channel or aggregate port). Similarly configured ports are grouped based on hardware, administrative, and port parameter constraints. For example, LACP groups the ports with the same speed, duplex mode, native VLAN, VLAN range, and trunking status and type. After grouping the links into an EtherChannel, LACP adds the group to the spanning tree as a single switch port.

LACP Modes

Table 27-2 shows the user-configurable EtherChannel LACP modes for the **channel-group** interface configuration command.

Table 27-2 EtherChannel LACP Modes

Mode	Description
active	Places a port into an active negotiating state in which the port starts negotiations with other ports by sending LACP packets.
passive	Places a port into a passive negotiating state in which the port responds to LACP packets that it receives, but does not start LACP packet negotiation. This setting minimizes the transmission of LACP packets.
on	Forces a port to channel without PAgP or LACP. With the on mode, a usable EtherChannel exists only when a port group in the on mode is connected to another port group in the on mode.

Both the **active** and **passive** LACP modes enable ports to negotiate with partner ports to form an EtherChannel based on criteria such as port speed and, for Layer 2 EtherChannels, trunking state and VLAN numbers.

Ports can form an EtherChannel when they are in different LACP modes as long as the modes are compatible. For example:

- A port in the **active** mode can form an EtherChannel with another port that is in the **active** or **passive** mode.
- A port in the **passive** mode cannot form an EtherChannel with another port that is also in the **passive** mode because neither port starts LACP negotiation.

LACP Interaction with Other Features

The DTP and the CDP send and receive packets over the physical ports in the EtherChannel. Trunk ports send and receive LACP PDUs on the lowest numbered VLAN.

In Layer 2 EtherChannels, the first port in the channel that comes up provides its MAC address to the EtherChannel. If this port is removed from the bundle, one of the remaining ports in the bundle provides its MAC address to the EtherChannel.

LACP sends and receives LACP PDUs only from ports that are up and have LACP enabled for the active or passive mode.

Load Balancing and Forwarding Methods

EtherChannel balances the traffic load across the links in a channel by reducing part of the binary pattern formed from the addresses in the frame to a numerical value that selects one of the links in the channel. EtherChannel load balancing can use MAC addresses or IP addresses, source or destination addresses, or both source and destination addresses. The selected mode applies to all EtherChannels configured on the switch. You configure the load balancing and forwarding method by using the **port-channel load-balance** global configuration command.

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.

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.

With source-and-destination MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on both the source and destination MAC addresses. This forwarding method, a combination source-MAC and destination-MAC address forwarding methods of load distribution, can be used if it is not clear whether source-MAC or destination-MAC address forwarding is better suited on a particular switch. With source-and-destination MAC-address forwarding, packets sent from host A to host B, host A to host C, and host C to host B could all use different ports in the channel.


With source-IP address-based forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the EtherChannel based on the source-IP address of the incoming packet. Therefore, to provide load-balancing, packets from different IP addresses use different ports in the channel, but packets from the same IP address use the same port in the channel.

With destination-IP address-based forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the EtherChannel based on the destination-IP address of the incoming packet. Therefore, to provide load-balancing, packets from the same IP source address sent to different IP destination addresses could be sent on different ports in the channel. But packets sent from different source IP addresses to the same destination IP address are always sent on the same port in the channel.

With source-and-destination IP address-based forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the EtherChannel based on both the source and destination IP addresses of the incoming packet. This forwarding method, a combination of source-IP and destination-IP address-based forwarding, can be used if it is not clear whether source-IP or destination-IP address-based forwarding is better suited on a particular switch. In this method, packets sent from the IP address A to IP address B, from IP address A to IP address C, and from IP address C to IP address B could all use different ports in the channel.

Different load-balancing methods have different advantages, and the choice of a particular load-balancing method should be based on the position of the switch in the network and the kind of traffic that needs to be load-distributed. In [Figure 27-3](#), an EtherChannel of four workstations communicates with a router. Because the router is a single-MAC-address device, source-based forwarding on the switch EtherChannel ensures that the switch uses all available bandwidth to the router. The router is configured for destination-based forwarding because the large number of workstations ensures that the traffic is evenly distributed from the router 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 27-3 Load Distribution and Forwarding Methods

Configuring EtherChannels

These sections describe how to configure EtherChannel on Layer 2 and Layer 3 ports:

- [Default EtherChannel Configuration, page 27-9](#)
- [EtherChannel Configuration Guidelines, page 27-9](#)
- [Configuring Layer 2 EtherChannels, page 27-10](#) (required)
- [Configuring Layer 3 EtherChannels, page 27-12](#) (required)
- [Configuring EtherChannel Load Balancing, page 27-15](#) (optional)
- [Configuring the PAGP Learn Method and Priority, page 27-16](#) (optional)
- [Configuring LACP Hot-Standby Ports, page 27-17](#) (optional)



Note

Make sure that the ports are correctly configured. For more information, see the [“EtherChannel Configuration Guidelines”](#) section on page 27-9.



Note

After you configure an EtherChannel, configuration changes applied to the port-channel interface apply to all the physical ports assigned to the port-channel interface, and configuration changes applied to the physical port affect only the port where you apply the configuration.

Default EtherChannel Configuration

Table 27-3 shows the default EtherChannel configuration.

Table 27-3 Default EtherChannel Configuration

Feature	Default Setting
Channel groups	None assigned.
Port-channel logical interface	None defined.
PAgP mode	No default.
PAgP learn method	Aggregate-port learning on all ports.
PAgP priority	128 on all ports.
LACP mode	No default.
LACP learn method	Aggregate-port learning on all ports.
LACP port priority	32768 on all ports.
LACP system priority	32768.
LACP system ID	LACP system priority and the switch MAC address.
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 ports are automatically disabled to avoid network loops and other problems. Follow these guidelines to avoid configuration problems:

- More than 12 EtherChannels cannot be configured on the switch.
- Do not bundle standard Gigabit Ethernet ports with the enhanced-services (ES) Gigabit Ethernet ports.
- Configure a PAgP EtherChannel with up to eight Ethernet ports of the same type.
- Configure a LACP EtherChannel with up to 16 Ethernet ports of the same type. Up to eight ports can be active, and up to eight ports can be in standby mode.
- Configure all ports in an EtherChannel to operate at the same speeds and duplex modes.
- Enable all ports in an EtherChannel. A port 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 ports 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 port to be a member of more than one EtherChannel group.

- Do not configure an EtherChannel in both the PAgP and LACP modes. EtherChannel groups running PAgP and LACP can coexist on the same switch. Individual EtherChannel groups can run either PAgP or LACP, but they cannot interoperate.
- Do not configure a Switched Port Analyzer (SPAN) destination as part of an EtherChannel.
- Do not configure a secure port as part of an EtherChannel or the reverse.
- Do not configure a port that is an active member of an EtherChannel as an 802.1x port. If 802.1x is enabled on a not-yet active port of an EtherChannel, the port does not join the EtherChannel.
- For Layer 2 EtherChannels:
 - Assign all ports in the EtherChannel to the same VLAN, or configure them as trunks. Ports with different native VLANs cannot form an EtherChannel.
 - If you configure an EtherChannel from trunk ports, verify that the trunking mode (ISL or 802.1Q) is the same on all the trunks. Inconsistent trunk modes on EtherChannel ports can have unexpected results.
 - An EtherChannel supports the same allowed range of VLANs on all the ports in a trunking Layer 2 EtherChannel. If the allowed range of VLANs is not the same, the ports do not form an EtherChannel even when PAgP is set to the **auto** or **desirable** mode.
 - Ports 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 ports incompatible for the formation of an EtherChannel.
- For Layer 3 EtherChannels, assign the Layer 3 address to the port-channel logical interface, not to the physical ports in the channel.

Configuring Layer 2 EtherChannels

You configure Layer 2 EtherChannels by assigning ports to a channel group with the **channel-group** interface configuration command. This command automatically creates the port-channel logical interface.

Beginning in privileged EXEC mode, follow these steps to assign a Layer 2 Ethernet port to a Layer 2 EtherChannel. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify a physical port to configure, and enter interface configuration mode. Valid interfaces include physical ports. For a PAgP EtherChannel, you can configure up to eight ports of the same type and speed for the same group. For a LACP EtherChannel, you can configure up to 16 Ethernet ports of the same type. Up to eight ports can be active, and up to eight ports can be in standby mode.
Step 3	switchport mode { access trunk } switchport access vlan <i>vlan-id</i>	Assign all ports as static-access ports in the same VLAN, or configure them as trunks. If you configure the port as a static-access port, assign it to only one VLAN. The range is 1 to 4094.

	Command	Purpose
Step 4	channel-group <i>channel-group-number</i> mode { auto [non-silent] desirable [non-silent] on } { active passive }	Assign the port to a channel group, and specify the PAgP or the LACP mode. For <i>channel-group-number</i>, the range is 1 to 12. For mode, select one of these keywords: • auto—Enables PAgP only if a PAgP device is detected. It places the port into a passive negotiating state, in which the port responds to PAgP packets it receives but does not start PAgP packet negotiation. • desirable—Unconditionally enables PAgP. It places the port into an active negotiating state, in which the port starts negotiations with other ports by sending PAgP packets. • on—Forces the port to channel without PAgP or LACP. With the on mode, a usable EtherChannel exists only when a port group in the on mode is connected to another port group in the on mode. • non-silent—(Optional) If your switch is connected to a partner that is PAgP-capable, configure the switch port for nonsilent operation when the port is in the auto or desirable mode. If you do not specify non-silent, silent is assumed. The silent setting is for connections to file servers or packet analyzers. This setting allows PAgP to operate, to attach the port to a channel group, and to use the port for transmission. • active—Enables LACP only if a LACP device is detected. It places the port into an active negotiating state in which the port starts negotiations with other ports by sending LACP packets. • passive—Enables LACP on the port and places it into a passive negotiating state in which the port responds to LACP packets that it receives, but does not start LACP packet negotiation. For information on compatible modes for the switch and its partner, see the “PAgP Modes” section on page 27-4 and the “LACP Modes” section on page 27-6.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a port from the EtherChannel group, use the **no channel-group** interface configuration command.

This example shows how to configure an EtherChannel. It assigns two ports as static-access ports in VLAN 10 to channel 5 with the PAgP mode **desirable**:

```
Switch# configure terminal
Switch(config)# interface range fastethernet1/0/4 -5
Switch(config-if-range)# switchport mode access
Switch(config-if-range)# switchport access vlan 10
Switch(config-if-range)# channel-group 5 mode desirable non-silent
Switch(config-if-range)# end
```

This example shows how to configure an EtherChannel. It assigns two ports as static-access ports in VLAN 10 to channel 5 with the LACP mode **active**:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet1/0/1 -2
Switch(config-if-range)# switchport mode access
Switch(config-if-range)# switchport access vlan 10
Switch(config-if-range)# channel-group 5 mode active
Switch(config-if-range)# end
```

Configuring Layer 3 EtherChannels

To configure Layer 3 EtherChannels, you create the port-channel logical interface and then put the Ethernet ports into the port-channel as described in the next two sections.

Creating Port-Channel Logical Interfaces

When configuring Layer 3 EtherChannels, you should first manually create the port-channel logical interface by using the **interface port-channel** global configuration command. Then you put the logical interface into the channel group by using the **channel-group** interface configuration command.



Note

To move an IP address from a physical port to an EtherChannel, you must delete the IP address from the physical port before configuring it on the port-channel interface. Beginning in privileged EXEC mode, follow these steps to create a port-channel interface for a Layer 3 EtherChannel. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface port-channel <i>port-channel-number</i>	Create the port-channel logical interface, and enter interface configuration mode. For <i>port-channel-number</i> , the range is 1 to 12.
Step 3	no switchport	Put the interface into Layer 3 mode.
Step 4	ip address <i>ip-address mask</i>	Assign an IP address and subnet mask to the EtherChannel.
Step 5	end	Return to privileged EXEC mode.
Step 6	show etherchannel <i>channel-group-number detail</i>	Verify your entries.

	Command	Purpose
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.
Step 8		Assign an Ethernet port to the Layer 3 EtherChannel. For more information, see the “Configuring the Physical Interfaces” section on page 27-13.

To remove the port-channel, use the **no interface port-channel** *port-channel-number* global configuration command.

This example shows how to create logical port-channel 5 and assign 172.10.20.10 as its IP address:

```
Switch# configure terminal
Switch(config)# interface port-channel 5
Switch(config-if)# no switchport
Switch(config-if)# ip address 172.10.20.10 255.255.255.0
Switch(config-if)# end
```

Configuring the Physical Interfaces

Beginning in privileged EXEC mode, follow these steps to assign an Ethernet port to a Layer 3 EtherChannel. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify a physical port to configure, and enter interface configuration mode. For a PAgP EtherChannel, you can configure up to eight ports of the same type and speed for the same group. For a LACP EtherChannel, you can configure up to 16 Ethernet ports of the same type. Up to eight ports can be active, and up to eight ports can be in standby mode.
Step 3	no ip address	Ensure that there is no IP address assigned to the physical port.
Step 4	no switchport	Put the port into Layer 3 mode.

	Command	Purpose
Step 5	channel-group <i>channel-group-number</i> mode { auto [non-silent] desirable [non-silent] on } { active passive }	Assign the port to a channel group, and specify the PAgP or the LACP mode. For <i>channel-group-number</i>, the range is 1 to 12. This number must be the same as the <i>port-channel-number</i> (logical port) configured in the “Creating Port-Channel Logical Interfaces” section on page 27-12. For mode, select one of these keywords: • auto—Enables PAgP only if a PAgP device is detected. It places the port into a passive negotiating state, in which the port responds to PAgP packets it receives but does not start PAgP packet negotiation. • desirable—Unconditionally enables PAgP. It places the port into an active negotiating state, in which the port starts negotiations with other ports by sending PAgP packets. • on—Forces the port to channel without PAgP or LACP. With the on mode, a usable EtherChannel exists only when a port group in the on mode is connected to another port group in the on mode. • non-silent—(Optional) If your switch is connected to a partner that is PAgP capable, configure the switch port for nonsilent operation when the port is in the auto or desirable mode. If you do not specify non-silent, silent is assumed. The silent setting is for connections to file servers or packet analyzers. This setting allows PAgP to operate, to attach the port to a channel group, and to use the port for transmission. • active—Enables LACP only if a LACP device is detected. It places the port into an active negotiating state in which the port starts negotiations with other ports by sending LACP packets. • passive—Enables LACP on the port and places it into a passive negotiating state in which the port responds to LACP packets that it receives, but does not start LACP packet negotiation. For information on compatible modes for the switch and its partner, see the “PAgP Modes” section on page 27-4 and the “LACP Modes” section on page 27-6.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to configure an EtherChannel. It assigns two ports to channel 5 with the LACP mode **active**:

```
Switch# configure terminal
Switch(config)# interface range fastethernet1/0/4 -5
Switch(config-if-range)# no ip address
Switch(config-if-range)# no switchport
Switch(config-if-range)# channel-group 5 mode active
Switch(config-if-range)# 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 [“Load Balancing and Forwarding Methods” section on page 27-6](#).

Beginning in privileged EXEC mode, follow these steps to configure EtherChannel load balancing. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	port-channel load-balance {dst-ip dst-mac src-dst-ip src-dst-mac src-ip src-mac}	Configure an EtherChannel load-balancing method. The default is src-mac. Select one of these load-distribution methods: • dst-ip—Load distribution is based on the destination-host IP address. • dst-mac—Load distribution is based on the destination-host MAC address of the incoming packet. • src-dst-ip—Load distribution is based on the source-and-destination host-IP address. • src-dst-mac—Load distribution is based on the source-and-destination host-MAC address. • src-ip—Load distribution is based on the source-host IP address. • src-mac—Load distribution is based on the source-MAC address of the incoming packet.
Step 3	end	Return to privileged EXEC mode.
Step 4	show etherchannel load-balance	Verify your entries.
Step 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 (logical) ports. The learn method must be configured the same at both ends of the link.

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 ports in the EtherChannel. With aggregate-port learning, it is not important on which physical port the packet arrives.

PAgP cannot automatically detect when the partner device is a physical learner and when the local device is an aggregate-port learner. Therefore, you must manually set the learning method on the local device to learn addresses by physical ports. You also must set the load-distribution method to source-based distribution, so that any given source MAC address is always sent on the same physical port.

You also can configure a single port within the group for all transmissions and use other ports for hot standby. The unused ports in the group can be swapped into operation in just a few seconds if the selected single port loses hardware-signal detection. You can configure which port is always selected for packet transmission by changing its priority with the **pagp port-priority** interface configuration command. The higher the priority, the more likely that the port will be selected.



Note

The switch supports address learning only on aggregate ports even though the **physical-port** keyword is provided in the CLI. The **pagp learn-method** command and the **pagp port-priority** command have no effect on the switch hardware, but they are required for PAgP interoperability with devices that only support address learning by physical ports, such as the Catalyst 1900 switch.

When the link partner to the switch is a physical learner (such as a Catalyst 1900 series switch), we recommend that you configure the switch as a physical-port learner by using the **pagp learn-method physical-port** interface configuration command. Set the load-distribution method based on the source MAC address by using the **port-channel load-balance src-mac** global configuration command. The switch then sends packets to the Catalyst 1900 switch using the same port in the EtherChannel from which it learned the source address. Use the **pagp learn-method** command only in this situation.

Beginning in privileged EXEC mode, follow these steps to configure your switch as a PAgP physical-port learner and to adjust the priority so that the same port in the bundle is selected for sending packets. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port for transmission, and enter interface configuration mode.

	Command	Purpose
Step 3	pagp learn-method physical-port	Select the PAgP learning method. By default, aggregation-port learning is selected, which means the switch sends packets to the source by using any of the ports in the EtherChannel. With aggregate-port learning, it is not important on which physical port the packet arrives. Select physical-port to connect with another switch that is a physical learner. Make sure to configure the port-channel load-balance global configuration command to src-mac as described in the “Configuring EtherChannel Load Balancing” section on page 27-15. The learning method must be configured the same at both ends of the link.
Step 4	pagp port-priority priority	Assign a priority so that the selected port is chosen for packet transmission. For <i>priority</i>, the range is 0 to 255. The default is 128. The higher the priority, the more likely that the port will be used for PAgP transmission.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config or show pagp channel-group-number internal	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the priority to its default setting, use the **no pagp port-priority** interface configuration command. To return the learning method to its default setting, use the **no pagp learn-method** interface configuration command.

Configuring LACP 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. The software places any additional links in a hot-standby mode. If one of the active links becomes inactive, a link that is in the hot-standby mode becomes active in its place.

If you configure more than eight links for an EtherChannel group, the software automatically decides which of the hot-standby ports to make active based on the LACP priority. The software assigns to every link between systems that operate LACP a unique priority made up of these elements (in priority order):

- LACP system priority
- System ID (a combination of the LACP system priority and the switch MAC address)
- LACP port priority
- Port number

In priority comparisons, numerically lower values have higher priority. The priority decides which ports should be put in standby mode when there is a hardware limitation that prevents all compatible ports from aggregating.

Ports are considered for active use in aggregation in link-priority order starting with the port attached to the highest priority link. Each port is selected for active use if the preceding higher priority selections can also be maintained. Otherwise, the port is selected for standby mode.

You can change the default values of the LACP system priority and the LACP port priority to affect how the software selects active and standby links. For more information, see the [“Configuring the LACP System Priority” section on page 27-18](#) and the [“Configuring the LACP Port Priority” section on page 27-18](#).

Configuring the LACP System Priority

You can configure the system priority for all of the EtherChannels that are enabled for LACP by using the **lacp system-priority** global configuration command. You cannot configure a system priority for each LACP-configured channel. By changing this value from the default, you can affect how the software selects active and standby links.

You can use the **show etherchannel summary** privileged EXEC command to see which ports are in the hot-standby mode (denoted with an *H* port-state flag).

Beginning in privileged EXEC mode, follow these steps to configure the LACP system priority. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	lacp system-priority <i>priority</i>	Configure the LACP system priority. For <i>priority</i> , the range is 1 to 65535. The default is 32768. The lower the value, the higher the system priority.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config or show lacp sys-id	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the LACP system priority to the default value, use the **no lacp system-priority** global configuration command.

Configuring the LACP Port Priority

By default, all ports use the same port priority. If the local system has a lower value for the system priority and the system ID than the remote system, you can affect which of the hot-standby links become active first by changing the port priority of LACP EtherChannel ports to a lower value than the default. The hot-standby ports that have lower port numbers become active in the channel first. You can use the **show etherchannel summary** privileged EXEC command to see which ports are in the hot-standby mode (denoted with an *H* port-state flag).



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 the hot-standby state and are used only if one of the channeled ports fails.

Beginning in privileged EXEC mode, follow these steps to configure the LACP port priority. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port to be configured, and enter interface configuration mode.
Step 3	lacp port-priority <i>priority</i>	Configure the LACP port priority. For <i>priority</i> , the range is 1 to 65535. The is 32768. The lower the value, the more likely that the port will be used for LACP transmission.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config or show lacp [<i>channel-group-number</i>] internal	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the LACP port priority to the default value, use the **no lacp port-priority** interface configuration command.

Displaying EtherChannel, PAgP, and LACP Status

To display EtherChannel, PAgP, and LACP status information, use the privileged EXEC commands described in [Table 27-4](#):

Table 27-4 Commands for Displaying EtherChannel, PAgP, and LACP Status

Command	Description
show etherchannel [<i>channel-group-number</i> { detail port port-channel protocol summary }] { detail load-balance port port-channel protocol summary }	Displays EtherChannel information in a brief, detailed, and one-line summary form. Also displays the load-balance or frame-distribution scheme, port, port-channel, and protocol information.
show pagp [<i>channel-group-number</i>] { counters internal neighbor }	Displays PAgP information such as traffic information, the internal PAgP configuration, and neighbor information.
show lacp [<i>channel-group-number</i>] { counters internal neighbor }	Displays LACP information such as traffic information, the internal LACP configuration, and neighbor information.

You can clear PAgP channel-group information and traffic counters by using the **clear pagp** {*channel-group-number* **counters** | **counters**} privileged EXEC command.

You can clear LACP channel-group information and traffic counters by using the **clear lacp** {*channel-group-number* **counters** | **counters**} privileged EXEC command.

For detailed information about the fields in the displays, refer to the command reference for this release.



Configuring IP Unicast Routing

This chapter describes how to configure IP unicast routing on the Catalyst 3750 Metro switch.



Note

For more detailed IP unicast configuration information, refer to the *Cisco IOS IP and IP Routing Configuration Guide for Release 12.1*. For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

This chapter consists of these sections:

- [Understanding IP Routing, page 28-2](#)
- [Steps for Configuring Routing, page 28-3](#)
- [Configuring IP Addressing, page 28-4](#)
- [Enabling IP Unicast Routing, page 28-17](#)
- [Configuring RIP, page 28-18](#)
- [Configuring IGRP, page 28-23](#)
- [Configuring OSPF, page 28-28](#)
- [Configuring EIGRP, page 28-37](#)
- [Configuring BGP, page 28-43](#)
- [Configuring ISO CLNS Routing, page 28-63](#)
- [Configuring Multi-VRF CE, page 28-76](#)
- [Configuring Protocol-Independent Features, page 28-86](#)
- [Monitoring and Maintaining the IP Network, page 28-100](#)



Note


When configuring routing parameters on the switch and to allocate system resources to maximize the number of unicast routes allowed, you can use the **sdm prefer routing** global configuration command to set the Switch Database Management (SDM) feature to the routing template. For more information on the SDM templates, see [Chapter 6, “Configuring SDM Templates,”](#) or refer to the **sdm prefer** command in the command reference for this release.

Understanding IP Routing

In some network environments, VLANs are associated with individual networks or subnetworks. In an IP network, each subnetwork is mapped to an individual VLAN. Configuring VLANs helps control the size of the broadcast domain and keeps local traffic local. However, network devices in different VLANs cannot communicate with one another without a Layer 3 device (router) to route traffic between the VLAN, referred to as inter-VLAN routing. You configure one or more routers to route traffic to the appropriate destination VLAN.

Figure 28-1 shows a basic routing topology. Switch A is in VLAN 10, and Switch B is in VLAN 20. The router has an interface in each VLAN.

Figure 28-1 Routing Topology Example



When Host A in VLAN 10 needs to communicate with Host B in VLAN 10, it sends a packet addressed to that host. Switch A forwards the packet directly to Host B, without sending it to the router.

When Host A sends a packet to Host C in VLAN 20, Switch A forwards the packet to the router, which receives the traffic on the VLAN 10 interface. The router checks the routing table, determines the correct outgoing interface, and forwards the packet on the VLAN 20 interface to Switch B. Switch B receives the packet and forwards it to Host C.

Routers and Layer 3 switches can route packets in three different ways:

- By using default routing
- By using preprogrammed static routes for the traffic
- By dynamically calculating routes by using a routing protocol

Default routing refers to sending traffic with a destination unknown to the router to a default outlet or destination.

Static unicast routing forwards packets from predetermined ports through a single path into and out of a network. Static routing is secure and uses little bandwidth, but does not automatically respond to changes in the network, such as link failures, and therefore, might result in unreachable destinations. As networks grow, static routing becomes a labor-intensive liability.

Dynamic routing protocols are used by routers to dynamically calculate the best route for forwarding traffic. There are two types of dynamic routing protocols:

- Routers using distance-vector protocols maintain routing tables with distance values of networked resources, and periodically pass these tables to their neighbors. Distance-vector protocols use one or a series of metrics for calculating the best routes. These protocols are easy to configure and use.
- Routers using link-state protocols maintain a complex database of network topology, based on the exchange of link-state advertisements (LSAs) between routers. LSAs are triggered by an event in the network, which speeds up the convergence time or time required to respond to these changes. Link-state protocols respond quickly to topology changes, but require greater bandwidth and more resources than distance-vector protocols.

Distance-vector protocols supported by the Catalyst 3750 switch are Routing Information Protocol (RIP), which uses a single distance metric (cost) to determine the best path; Interior Gateway Routing Protocol (IGRP), which uses a series of metrics; and Border Gateway Protocol (BGP), which adds a path vector mechanism. The switch also supports the Open Shortest Path First (OSPF) link-state protocol and Enhanced IGRP (EIGRP), which adds some link-state routing features to traditional IGRP to improve efficiency. The switch supports ISO connectionless dynamic routing protocols, such as Interior Gateway Routing Protocol (IGRP), which is a distance vector protocol, and Intermediate System-to-Intermediate System (IS-IS) routing.

Steps for Configuring Routing

By default, IP routing is disabled on the switch, and you must enable it before routing can take place. For detailed IP routing configuration information, refer to the *Cisco IOS IP and IP Routing Configuration Guide for Release 12.1*.

In the following procedures, the specified interface must be one of these Layer 3 interfaces:

- A routed port: a physical port configured as a Layer 3 port by using the **no switchport** interface configuration command.
- A switch virtual interface (SVI): a VLAN interface created by using the **interface vlan** *vlan_id* global configuration command and by default a Layer 3 interface.
- An EtherChannel port channel in Layer 3 mode: a port-channel logical interface created by using the **interface port-channel** *port-channel-number* global configuration command and binding the Ethernet interface into the channel group. For more information, see the [“Configuring Layer 3 EtherChannels” section on page 27-12](#).



Note

A Layer 3 switch can have an IP address assigned to each routed port and SVI. The number of routed ports and SVIs that you can configure is not limited by software. However, the interrelationship between this number and the number and volume of features being implemented might have an impact on CPU utilization because of hardware limitations. To optimize system memory for routing, use the **sdm prefer routing** global configuration command.

All Layer 3 interfaces on which routing will occur must have IP addresses assigned to them. See the [“Assigning IP Addresses to Network Interfaces” section on page 28-5](#).

Configuring routing consists of several main procedures:

- To support VLAN interfaces, create and configure VLANs on the switch, and assign VLAN membership to Layer 2 interfaces. For more information, see [Chapter 10, “Configuring VLANs.”](#)
- Configure Layer 3 interfaces.
- Enable IP routing on the switch.
- Assign IP addresses to the Layer 3 interfaces.
- Enable selected routing protocols on the switch.
- Configure routing protocol parameters (optional).

Configuring IP Addressing

A required task for configuring IP routing is to assign IP addresses to Layer 3 network interfaces to enable the interfaces and allow communication with the hosts on those interfaces that use IP. These sections describe how to configure various IP addressing features. Assigning IP addresses to the interface is required; the other procedures are optional.

- [Default Addressing Configuration, page 28-4](#)
- [Assigning IP Addresses to Network Interfaces, page 28-5](#)
- [Configuring Address Resolution Methods, page 28-7](#)
- [Routing Assistance When IP Routing is Disabled, page 28-10](#)
- [Configuring Broadcast Packet Handling, page 28-12](#)
- [Monitoring and Maintaining IP Addressing, page 28-16](#)

Default Addressing Configuration

[Table 28-1](#) shows the default addressing configuration.

Table 28-1 Default Addressing Configuration

Feature	Default Setting
IP address	None defined.
ARP	No permanent entries in the Address Resolution Protocol (ARP) cache. Encapsulation: Standard Ethernet-style ARP. Timeout: 14400 seconds (4 hours).
IP broadcast address	255.255.255.255 (all ones).
IP classless routing	Enabled.
IP default gateway	Disabled.
IP directed broadcast	Disabled (all IP directed broadcasts are dropped).
IP domain	Domain list: No domain names defined. Domain lookup: Enabled. Domain name: Enabled.
IP forward-protocol	If a helper address is defined or User Datagram Protocol (UDP) flooding is configured, UDP forwarding is enabled on default ports. Any-local-broadcast: Disabled. Spanning Tree Protocol (STP): Disabled. Turbo-flood: Disabled.
IP helper address	Disabled.
IP host	Disabled.

Table 28-1 Default Addressing Configuration (continued)

Feature	Default Setting
IRDP	Disabled. Defaults when enabled: - Broadcast IRDP advertisements. - Maximum interval between advertisements: 600 seconds. - Minimum interval between advertisements: 0.75 times max interval - Preference: 0.
IP proxy ARP	Enabled.
IP routing	Disabled.
IP subnet-zero	Disabled.

Assigning IP Addresses to Network Interfaces

An IP address identifies a location to which IP packets can be sent. Some IP addresses are reserved for special uses and cannot be used for host, subnet, or network addresses. RFC 1166, “Internet Numbers,” contains the official description of IP addresses.

An interface can have one primary IP address. A mask identifies the bits that denote the network number in an IP address. When you use the mask to subnet a network, the mask is referred to as a subnet mask. To receive an assigned network number, contact your Internet service provider.

Beginning in privileged EXEC mode, follow these steps to assign an IP address and a network mask to a Layer 3 interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	no switchport	Remove the interface from Layer 2 configuration mode (if it is a physical interface).
Step 4	ip address <i>ip-address subnet-mask</i>	Configure the IP address and IP subnet mask.
Step 5	no shutdown	Enable the interface.
Step 6	end	Return to privileged EXEC mode.
Step 7	show interfaces [<i>interface-id</i>] show ip interface [<i>interface-id</i>] show running-config interface [<i>interface-id</i>]	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use of Subnet Zero

Subnetting with a subnet address of zero is strongly discouraged because of the problems that can arise if a network and a subnet have the same addresses. For example, if network 131.108.0.0 is subnetted as 255.255.255.0, subnet zero would be written as 131.108.0.0, which is the same as the network address.

You can use the all ones subnet (131.108.255.0) and even though it is discouraged, you can enable the use of subnet zero if you need the entire subnet space for your IP address.

Beginning in privileged EXEC mode, follow these steps to enable subnet zero:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip subnet-zero	Enable the use of subnet zero for interface addresses and routing updates.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.


Use the **no ip subnet-zero** global configuration command to restore the default and disable the use of subnet zero.

Classless Routing

By default, classless routing behavior is enabled on the switch when it is configured to route. With classless routing, if a router receives packets for a subnet of a network with no default route, the router forwards the packet to the best supernet route. A *supernet* consists of contiguous blocks of Class C address spaces used to simulate a single, larger address space and is designed to relieve the pressure on the rapidly depleting Class B address space.


In [Figure 28-2](#), classless routing is enabled. When the host sends a packet to 120.20.4.1, instead of discarding the packet, the router forwards it to the best supernet route. If you disable classless routing and a router receives packets destined for a subnet of a network with no network default route, the router discards the packet.

Figure 28-2 IP Classless Routing



In [Figure 28-3](#), the router in network 128.20.0.0 is connected to subnets 128.20.1.0, 128.20.2.0, and 128.20.3.0. If the host sends a packet to 120.20.4.1, because there is no network default route, the router discards the packet.

Figure 28-3 No IP Classless Routing



To prevent the switch from forwarding packets destined for unrecognized subnets to the best supernet route possible, you can disable classless routing behavior.

Beginning in privileged EXEC mode, follow these steps to disable classless routing:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no ip classless	Disable classless routing behavior.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To restore the default and have the switch forward packets destined for a subnet of a network with no network default route to the best supernet route possible, use the **ip classless** global configuration command.

Configuring Address Resolution Methods

You can control interface-specific handling of IP by using address resolution. A device using IP can have both a local address or MAC address, which uniquely defines the device on its local segment or LAN, and a network address, which identifies the network to which the device belongs.

The local address or MAC address is known as a data link address because it is contained in the data link layer (Layer 2) section of the packet header and is read by data link (Layer 2) devices. To communicate with a device on Ethernet, the software must determine the MAC address of the device. The process of determining the MAC address from an IP address is called *address resolution*. The process of determining the IP address from the MAC address is called *reverse address resolution*.

The switch can use these forms of address resolution:

- Address Resolution Protocol (ARP) is used to associate IP address with MAC addresses. Taking an IP address as input, ARP determines the associated MAC address and then stores the IP address/MAC address association 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 or replies on IEEE 802 networks other than Ethernet is specified by the Subnetwork Access Protocol (SNAP).
- Proxy ARP helps hosts with no routing tables determine the MAC addresses of hosts on other networks or subnets. If the switch (router) receives an ARP request for a host that is not on the same interface as the ARP request sender, and if the router has all of its routes to the host through other interfaces, it generates a proxy ARP packet giving its own local data link address. The host that sent the ARP request then sends its packets to the router, which forwards them to the intended host.

Catalyst 3750 Metro switches also use the Reverse Address Resolution Protocol (RARP), which functions the same as ARP does, except that the RARP packets request an IP address instead of a local MAC address. Using RARP requires a RARP server on the same network segment as the router interface. Use the **ip rarp-server address** interface configuration command to identify the server.

For more information on RARP, refer to the *Cisco IOS Configuration Fundamentals Configuration Guide for Release 12.1*.

You can perform these tasks to configure address resolution:

- [Define a Static ARP Cache, page 28-8](#)
- [Set ARP Encapsulation, page 28-9](#)
- [Enable Proxy ARP, page 28-10](#)

Define a Static ARP Cache

ARP and other address resolution protocols provide dynamic mapping between IP addresses and MAC addresses. Because most hosts support dynamic address resolution, you usually do not need to specify static ARP cache entries. If you must define a static ARP cache entry, you can do so globally, which installs a permanent entry in the ARP cache that the switch uses to translate IP addresses into MAC addresses. Optionally, you can also specify that the switch respond to ARP requests as if it were the owner of the specified IP address. If you do not want the ARP entry to be permanent, you can specify a timeout period for the ARP entry.

Beginning in privileged EXEC mode, follow these steps to provide static mapping between IP addresses and MAC addresses:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	arp ip-address hardware-address type	Globally associate an IP address with a MAC (hardware) address in the ARP cache, and specify encapsulation type as one of these: • arpa—ARP encapsulation for Ethernet interfaces • snap—Subnetwork Address Protocol encapsulation for Token Ring and FDDI interfaces • sap—HP's ARP type

	Command	Purpose
Step 3	arp <i>ip-address hardware-address type</i> [<i>alias</i>]	(Optional) Specify that the switch respond to ARP requests as if it were the owner of the specified IP address.
Step 4	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to configure.
Step 5	arp timeout <i>seconds</i>	(Optional) Set the length of time an ARP cache entry will stay in the cache. The default is 14400 seconds (4 hours). The range is 0 to 2147483 seconds.
Step 6	end	Return to privileged EXEC mode.
Step 7	show interfaces [<i>interface-id</i>]	Verify the type of ARP and the timeout value used on all interfaces or a specific interface.
Step 8	show arp or show ip arp	View the contents of the ARP cache.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove an entry from the ARP cache, use the **no arp** *ip-address hardware-address type* global configuration command. To remove all nonstatic entries from the ARP cache, use the **clear arp-cache** privileged EXEC command.

Set ARP Encapsulation

By default, Ethernet ARP encapsulation (represented by the **arpa** keyword) is enabled on an IP interface. You can change the encapsulation methods to SNAP if required by your network.

Beginning in privileged EXEC mode, follow these steps to specify the ARP encapsulation type:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	arp { arpa snap }	Specify the ARP encapsulation method: • arpa—Address Resolution Protocol • snap—Subnetwork Address Protocol
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces [<i>interface-id</i>]	Verify ARP encapsulation configuration on all interfaces or the specified interface.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable an encapsulation type, use the **no arp arpa** or **no arp snap** interface configuration command.

Enable Proxy ARP

By default, the switch uses proxy ARP to help hosts determine MAC addresses of hosts on other networks or subnets.

Beginning in privileged EXEC mode, follow these steps to enable proxy ARP if it has been disabled:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	ip proxy-arp	Enable proxy ARP on the interface.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip interface [<i>interface-id</i>]	Verify the configuration on the interface or all interfaces.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable proxy ARP on the interface, use the **no ip proxy-arp** interface configuration command.

Routing Assistance When IP Routing is Disabled

These mechanisms allow the switch to learn about routes to other networks when it does not have IP routing enabled:

- [Proxy ARP, page 28-10](#)
- [Default Gateway, page 28-10](#)
- [ICMP Router Discovery Protocol \(IRDP\), page 28-11](#)

Proxy ARP

Proxy ARP, the most common method for learning about other routes, enables an Ethernet host with no routing information to communicate with hosts on other networks or subnets. The host assumes that all hosts are on the same local Ethernet and that they can use ARP to determine their MAC addresses. If a switch receives an ARP request for a host that is not on the same network as the sender, the switch evaluates whether it has the best route to that host. If it does, it sends an ARP reply packet with its own Ethernet MAC address, and the host that sent the request sends the packet to the switch, which forwards it to the intended host. Proxy ARP treats all networks as if they are local and performs ARP requests for every IP address.

Proxy ARP is enabled by default. To enable it after it has been disabled, see the [“Enable Proxy ARP” section on page 28-10](#). Proxy ARP works as long as other routers support it.

Default Gateway

Another method for locating routes is to define a default router or default gateway. All nonlocal packets are sent to this router, which either routes them appropriately or sends an IP Control Message Protocol (ICMP) redirect message back, defining which local router the host should use. The switch caches the redirect messages and forwards each packet as efficiently as possible. A limitation of this method is that there is no means of detecting when the default router has gone down or is unavailable.

Beginning in privileged EXEC mode, follow these steps to define a default gateway (router) when IP routing is disabled:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip default-gateway <i>ip-address</i>	Set up a default gateway (router).
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip redirects	Display the address of the default gateway router to verify the setting.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip default-gateway** global configuration command to disable this function.

ICMP Router Discovery Protocol (IRDP)

Router discovery allows the switch to dynamically learn about routes to other networks using IRDP. IRDP allows hosts to locate routers. When operating as a client, the switch generates router discovery packets. When operating as a host, the switch receives router discovery packets. The switch can also listen to Routing Information Protocol (RIP) and Interior Gateway Routing Protocol (IGRP) routing updates and use this information to infer locations of routers. The switch does not actually store the routing tables sent by routing devices; it merely keeps track of which systems are sending the data. The advantage of using IRDP is that it allows each router to specify both a priority and the time after which a device is assumed to be down if no further packets are received.

Each device discovered becomes a candidate for the default router, and a new highest-priority router is selected when a higher priority router is discovered, when the current default router is declared down, or when a TCP connection is about to time out because of excessive retransmissions.

The only required task for IRDP routing on an interface is to enable IRDP processing on that interface. When enabled, the default parameters apply. You can optionally change any of these parameters.

Beginning in privileged EXEC mode, follow these steps to enable and configure IRDP on an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	ip irdp	Enable IRDP processing on the interface.
Step 4	ip irdp multicast	(Optional) Send IRDP advertisements to the multicast address (224.0.0.1) instead of IP broadcasts. Note This command allows for compatibility with Sun Microsystems Solaris, which requires IRDP packets to be sent out as multicasts. Many implementations cannot receive these multicasts; ensure end-host ability before using this command.
Step 5	ip irdp holdtime <i>seconds</i>	(Optional) Set the IRDP period for which advertisements are valid. The default is three times the maxadvertinterval value. It must be greater than maxadvertinterval and cannot be greater than 9000 seconds. If you change the maxadvertinterval value, this value also changes.

	Command	Purpose
Step 6	ip irdp maxadvertinterval <i>seconds</i>	(Optional) Set the IRDP maximum interval between advertisements. The default is 600 seconds.
Step 7	ip irdp minadvertinterval <i>seconds</i>	(Optional) Set the IRDP minimum interval between advertisements. The default is 0.75 times the maxadvertinterval . If you change the maxadvertinterval , this value changes to the new default (0.75 of maxadvertinterval).
Step 8	ip irdp preference <i>number</i>	(Optional) Set a device IRDP preference level. The allowed range is -2^{31} to 2^{31} . The default is 0. A higher value increases the router preference level.
Step 9	ip irdp address <i>address</i> [<i>number</i>]	(Optional) Specify an IRDP address and preference to proxy-advertise.
Step 10	end	Return to privileged EXEC mode.
Step 11	show ip irdp	Verify settings by displaying IRDP values.
Step 12	copy running-config startup-config	(Optional) Save your entries in the configuration file.

If you change the **maxadvertinterval** value, the **holdtime** and **minadvertinterval** values also change, so it is important to first change the **maxadvertinterval** value, before manually changing either the **holdtime** or **minadvertinterval** values.

Use the **no ip irdp** interface configuration command to disable IRDP routing.

Configuring Broadcast Packet Handling

After configuring an IP interface address, you can enable routing and configure one or more routing protocols, or you can configure the way the switch responds to network broadcasts. A broadcast is a data packet destined for all hosts on a physical network. The switch supports two kinds of broadcasting:

- A directed broadcast packet is sent to a specific network or series of networks. A directed broadcast address includes the network or subnet fields.
- A flooded broadcast packet is sent to every network.



Note

You can also limit broadcast, unicast, and multicast traffic on Layer 2 interfaces by using the **storm-control** interface configuration command to set traffic suppression levels. For more information, see [Chapter 18, “Configuring Port-Based Traffic Control.”](#)

Routers provide some protection from broadcast storms by limiting their extent to the local cable. Bridges (including intelligent bridges), because they are Layer 2 devices, forward broadcasts to all network segments, thus propagating broadcast storms. The best solution to the broadcast storm problem is to use a single broadcast address scheme on a network. In most modern IP implementations, you can set the address to be used as the broadcast address. Many implementations, including the one in the Catalyst 3750 switch, support several addressing schemes for forwarding broadcast messages.

Perform the tasks in these sections to enable these schemes:

- [Enabling Directed Broadcast-to-Physical Broadcast Translation, page 28-13](#)
- [Forwarding UDP Broadcast Packets and Protocols, page 28-14](#)
- [Establishing an IP Broadcast Address, page 28-15](#)
- [Flooding IP Broadcasts, page 28-15](#)

Enabling Directed Broadcast-to-Physical Broadcast Translation

By default, IP directed broadcasts are dropped; they are not forwarded. Dropping IP-directed broadcasts makes routers less susceptible to denial-of-service attacks.

You can enable forwarding of IP-directed broadcasts on an interface where the broadcast becomes a physical (MAC-layer) broadcast. Only those protocols configured by using the **ip forward-protocol** global configuration command are forwarded.

You can specify an access list to control which broadcasts are forwarded. When an access list is specified, only those IP packets permitted by the access list are eligible to be translated from directed broadcasts to physical broadcasts. For more information on access lists, see [Chapter 25, “Configuring Network Security with ACLs.”](#)

Beginning in privileged EXEC mode, follow these steps to enable forwarding of IP-directed broadcasts on an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to configure.
Step 3	ip directed-broadcast [<i>access-list-number</i>]	Enable directed broadcast-to-physical broadcast translation on the interface. You can include an access list to control which broadcasts are forwarded. When an access list is specified, only IP packets permitted by the access list are eligible to be translated.
Step 4	exit	Return to global configuration mode.
Step 5	ip forward-protocol { udp [<i>port</i>] nd sdns }	Specify which protocols and ports the router forwards when forwarding broadcast packets. • udp—Forward UPD datagrams. <i>port</i>: (Optional) Destination port that controls which UDP services are forwarded. • nd—Forward ND datagrams. • sdns—Forward SDNS datagrams
Step 6	end	Return to privileged EXEC mode.
Step 7	show ip interface [<i>interface-id</i>] or show running-config	Verify the configuration on the interface or all interfaces.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip directed-broadcast** interface configuration command to disable translation of directed broadcast to physical broadcasts. Use the **no ip forward-protocol** global configuration command to remove a protocol or port.

Forwarding UDP Broadcast Packets and Protocols

User Datagram Protocol (UDP) is an IP host-to-host layer protocol, as is TCP. UDP provides a low-overhead, connectionless session between two end systems and does not provide for acknowledgment of received datagrams. Network hosts occasionally use UDP broadcasts to determine address, configuration, and name information. If such a host is on a network segment that does not include a server, UDP broadcasts are normally not forwarded. You can remedy this situation by configuring an interface on a router to forward certain classes of broadcasts to a helper address. You can use more than one helper address per interface.

You can specify a UDP destination port to control which UDP services are forwarded. You can specify multiple UDP protocols. You can also specify the Network Disk (ND) protocol, which is used by older diskless Sun workstations and the network security protocol SDNS.

By default, both UDP and ND forwarding are enabled if a helper address has been defined for an interface. The description for the **ip forward-protocol** interface configuration command in the *Cisco IOS IP and IP Routing Command Reference for Release 12.1* lists the ports that are forwarded by default if you do not specify any UDP ports.

If you do not specify any UDP ports when you configure the forwarding of UDP broadcasts, you are configuring the router to act as a BOOTP forwarding agent. BOOTP packets carry DHCP information.

Beginning in privileged EXEC mode, follow these steps to enable forwarding UDP broadcast packets on an interface and specify the destination address:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	ip helper-address <i>address</i>	Enable forwarding and specify the destination address for forwarding UDP broadcast packets, including BOOTP.
Step 4	exit	Return to global configuration mode.
Step 5	ip forward-protocol { udp [<i>port</i>] nd sdns }	Specify which protocols the router forwards when forwarding broadcast packets.
Step 6	end	Return to privileged EXEC mode.
Step 7	show ip interface [<i>interface-id</i>] or show running-config	Verify the configuration on the interface or all interfaces.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip helper-address** interface configuration command to disable the forwarding of broadcast packets to specific addresses. Use the **no ip forward-protocol** global configuration command to remove a protocol or port.

Establishing an IP Broadcast Address

The most popular IP broadcast address (and the default) is an address consisting of all ones (255.255.255.255). However, the switch can be configured to generate any form of IP broadcast address.

Beginning in privileged EXEC mode, follow these steps to set the IP broadcast address on an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to configure.
Step 3	ip broadcast-address <i>ip-address</i>	Enter a broadcast address different from the default, for example 128.1.255.255.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip interface [<i>interface-id</i>]	Verify the broadcast address on the interface or all interfaces.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To restore the default IP broadcast address, use the **no ip broadcast-address** interface configuration command.

Flooding IP Broadcasts

You can allow IP broadcasts to be flooded throughout your internetwork in a controlled fashion by using the database created by the bridging STP. Using this feature also prevents loops. To support this capability, bridging must be configured on each interface that is to participate in the flooding. If bridging is not configured on an interface, it still can receive broadcasts. However, the interface never forwards broadcasts it receives, and the router never uses that interface to send broadcasts received on a different interface.

Packets that are forwarded to a single network address using the IP helper-address mechanism can be flooded. Only one copy of the packet is sent on each network segment.

To be considered for flooding, packets must meet these criteria. (Note that these are the same conditions used to consider packet forwarding using IP helper addresses.)

- The packet must be a MAC-level broadcast.
- The packet must be an IP-level broadcast.
- The packet must be a TFTP, DNS, Time, NetBIOS, ND, or BOOTP packet, or a UDP specified by the **ip forward-protocol udp** global configuration command.
- The time-to-live (TTL) value of the packet must be at least two.

A flooded UDP datagram is given the destination address specified with the **ip broadcast-address** interface configuration command on the output interface. The destination address can be set to any address. Thus, the destination address might change as the datagram propagates through the network. The source address is never changed. The TTL value is decremented.

When a flooded UDP datagram is sent out an interface (and the destination address possibly changed), the datagram is handed to the normal IP output routines and is, therefore, subject to access lists, if they are present on the output interface.

Beginning in privileged EXEC mode, follow these steps to use the bridging spanning-tree database to flood UDP datagrams:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip forward-protocol spanning-tree	Use the bridging spanning-tree database to flood UDP datagrams.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

Use the **no ip forward-protocol spanning-tree** global configuration command to disable the flooding of IP broadcasts.

In a Catalyst 3750 Metro switch, the majority of packets are forwarded in hardware; most packets do not go through the switch CPU. For those packets that do go to the CPU, you can speed up spanning tree-based UDP flooding by a factor of about four to five times by using turbo-flooding. This feature is supported over Ethernet interfaces configured for ARP encapsulation.

Beginning in privileged EXEC mode, follow these steps to increase spanning-tree-based flooding:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode
Step 2	ip forward-protocol turbo-flood	Use the spanning-tree database to speed up flooding of UDP datagrams.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To disable this feature, use the **no ip forward-protocol turbo-flood** global configuration command.

Monitoring and Maintaining IP Addressing

When the contents of a particular cache, table, or database have become or are suspected to be invalid, you can remove all its contents by using the **clear** privileged EXEC commands. [Table 28-2](#) lists the commands for clearing contents.

Table 28-2 *Commands to Clear Caches, Tables, and Databases*

Command	Purpose
clear arp-cache	Clear the IP ARP cache and the fast-switching cache.
clear host {name *}	Remove one or all entries from the host name and the address cache.
clear ip route {network [mask] *}	Remove one or more routes from the IP routing table.

You can display specific statistics, such as the contents of IP routing tables, caches, and databases; the reachability of nodes; and the routing path that packets are taking through the network. [Table 28-3](#) lists the privileged EXEC commands for displaying IP statistics.

Table 28-3 Commands to Display Caches, Tables, and Databases

Command	Purpose
show arp	Display the entries in the ARP table.
show hosts	Display the default domain name, style of lookup service, name server hosts, and the cached list of host names and addresses.
show ip aliases	Display IP addresses mapped to TCP ports (aliases).
show ip arp	Display the IP ARP cache.
show ip interface <i>[interface-id]</i>	Display the IP status of interfaces.
show ip irdp	Display IRDP values.
show ip masks <i>address</i>	Display the masks used for network addresses and the number of subnets using each mask.
show ip redirects	Display the address of a default gateway.
show ip route <i>[address [mask]] [protocol]</i>	Display the current state of the routing table.
show ip route summary	Display the current state of the routing table in summary form.

Enabling IP Unicast Routing

By default, the switch is in Layer 2 switching mode and IP routing is disabled. To use the Layer 3 capabilities of the switch, you must enable IP routing.

Beginning in privileged EXEC mode, follow these steps to enable IP routing:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip routing	Enable IP routing.
Step 3	router <i>ip_routing_protocol</i>	Specify an IP routing protocol. This step might include other commands, such as specifying the networks to route with the network (RIP) router configuration command. For information on specific protocols, refer to sections later in this chapter and to the <i>Cisco IOS IP and IP Routing Configuration Guide for Release 12.1</i> .
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip routing** global configuration command to disable routing.

This example shows how to enable IP routing using RIP as the routing protocol:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip routing
Switch(config)# router rip
Switch(config-router)# network 10.0.0.0
Switch(config-router)# end
```

You can now set up parameters for the selected routing protocols as described in these sections:

- [Configuring RIP, page 28-18](#)
- [Configuring IGRP, page 28-23](#)
- [Configuring OSPF, page 28-28](#)
- [Configuring EIGRP, page 28-37](#)
- [Configuring BGP, page 28-43](#)
- [Configuring ISO CLNS Routing, page 28-63](#)
- [Configuring Multi-VRF CE, page 28-76](#)

You can also configure nonprotocol-specific features:

- [Configuring Protocol-Independent Features, page 28-86](#)

Configuring RIP

The Routing Information Protocol (RIP) is an interior gateway protocol (IGP) created for use in small, homogeneous networks. It is a distance-vector routing protocol that uses broadcast User Datagram Protocol (UDP) data packets to exchange routing information. The protocol is documented in RFC 1058. You can find detailed information about RIP in *IP Routing Fundamentals*, published by Cisco Press.

Using RIP, the switch sends routing information updates (advertisements) every 30 seconds. If a router does not receive an update from another router for 180 seconds or more, it marks the routes served by that router as unusable. If there is still no update after 240 seconds, the router removes all routing table entries for the non-updating router.

RIP uses hop counts to rate the value of different routes. The hop count is the number of routers that can be traversed in a route. A directly connected network has a hop count of zero; a network with a hop count of 16 is unreachable. This small range (0 to 15) makes RIP unsuitable for large networks.

If the router has a default network path, RIP advertises a route that links the router to the pseudonetwork 0.0.0.0. The 0.0.0.0 network does not exist; it is treated by RIP as a network to implement the default routing feature. The switch advertises the default network if a default was learned by RIP or if the router has a gateway of last resort and RIP is configured with a default metric. RIP sends updates to the interfaces in specified networks. If an interface's network is not specified, it is not advertised in any RIP update.

This section briefly describes how to configure RIP. It includes this information:

- [Default RIP Configuration, page 28-18](#)
- [Configuring Basic RIP Parameters, page 28-19](#)
- [Configuring RIP Authentication, page 28-21](#)
- [Configuring Summary Addresses and Split Horizon, page 28-21](#)

Default RIP Configuration

[Table 28-4](#) shows the default RIP configuration.

Table 28-4 Default RIP Configuration

Feature	Default Setting
Auto summary	Enabled.
Default-information originate	Disabled.
Default metric	Built-in; automatic metric translations.
IP RIP authentication key-chain	No authentication. Authentication mode: clear text.
IP RIP receive version	According to the version router configuration command.
IP RIP send version	According to the version router configuration command.
IP RIP triggered	According to the version router configuration command.
IP split horizon	Varies with media.
Neighbor	None defined.
Network	None specified.
Offset list	Disabled.
Output delay	0 milliseconds.
Timers basic	Update: 30 seconds. Invalid: 180 seconds. Hold-down: 180 seconds. Flush: 240 seconds.
Validate-update-source	Enabled.
Version	Receives RIP version 1 and 2 packets; sends version 1 packets.

Configuring Basic RIP Parameters

To configure RIP, you enable RIP routing for a network and optionally configure other parameters. Beginning in privileged EXEC mode, follow these steps to enable and configure RIP:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip routing	Enable IP routing. (Required only if IP routing is disabled.)
Step 3	router rip	Enable a RIP routing process, and enter router configuration mode.
Step 4	network <i>network number</i>	Associate a network with a RIP routing process. You can specify multiple network commands. RIP routing updates are sent and received through interfaces only on these networks.
Step 5	neighbor <i>ip-address</i>	(Optional) Define a neighboring router with which to exchange routing information. This step allows routing updates from RIP (normally a broadcast protocol) to reach nonbroadcast networks.

	Command	Purpose
Step 6	offset list [<i>access-list number</i> <i>name</i>] { <i>in</i> <i>out</i> } <i>offset</i> [<i>type number</i>]	(Optional) Apply an offset list to routing metrics to increase incoming and outgoing metrics to routes learned through RIP. You can limit the offset list with an access list or an interface.
Step 7	timers basic <i>update invalid holddown flush</i>	(Optional) Adjust routing protocol timers. Valid ranges for all timers are 0 to 4294967295 seconds. <i>update</i>—The time between sending routing updates. The default is 30 seconds. <i>invalid</i>—The timer after which a route is declared invalid. The default is 180 seconds. <i>holddown</i>—The time before a route is removed from the routing table. The default is 180 seconds. <i>flush</i>—The amount of time for which routing updates are postponed. The default is 240 seconds.
Step 8	version { <i>1</i> <i>2</i> }	(Optional) Configure the switch to receive and send only RIP Version 1 or RIP version 2 packets. By default, the switch receives Version 1 and 2 but sends only Version 1. You can also use the interface commands ip rip {send receive} version 1 2 1 2 to control what versions are used for sending and receiving on interfaces.
Step 9	no auto summary	(Optional) Disable automatic summarization. By default, the switch summarizes subprefixes when crossing classful network boundaries. Disable summarization (RIP version 2 only) to advertise subnet and host routing information to classful network boundaries.
Step 10	no validate-update-source	(Optional) Disable validation of the source IP address of incoming RIP routing updates. By default, the switch validates the source IP address of incoming RIP routing updates and discards the update if the source address is not valid. Under normal circumstances, disabling this feature is not recommended. However, if you have a router that is off-network and you want to receive its updates, you can use this command.
Step 11	output-delay <i>delay</i>	(Optional) Add interpacket delay for RIP updates sent. By default, packets in a multiple-packet RIP update have no delay added between packets. If you are sending packets to a lower-speed device, you can add an interpacket delay in the range of 8 to 50 milliseconds.
Step 12	end	Return to privileged EXEC mode.
Step 13	show ip protocols	Verify your entries.
Step 14	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To turn off the RIP routing process, use the **no router rip** global configuration command.

To display the parameters and current state of the active routing protocol process, use the **show ip protocols** privileged EXEC command. Use the **show ip rip database** privileged EXEC command to display summary address entries in the RIP database.

Configuring RIP Authentication

RIP version 1 does not support authentication. If you are sending and receiving RIP Version 2 packets, you can enable RIP authentication on an interface. The key chain determines the set of keys that can be used on the interface. If a key chain is not configured, no authentication is performed, not even the default. Therefore, you must also perform the tasks in the [“Managing Authentication Keys” section on page 28-99](#).

The switch supports two modes of authentication on interfaces for which RIP authentication is enabled: plain text and MD5. The default is plain text.

Beginning in privileged EXEC mode, follow these steps to configure RIP authentication on an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to configure.
Step 3	ip rip authentication key-chain <i>name-of-chain</i>	Enable RIP authentication.
Step 4	ip rip authentication mode [<i>text</i> <i>md5</i>]	Configure the interface to use plain text authentication (the default) or MD5 digest authentication.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config interface [<i>interface-id</i>]	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To restore clear text authentication, use the **no ip rip authentication mode** interface configuration command. To prevent authentication, use the **no ip rip authentication key-chain** interface configuration command.

Configuring Summary Addresses and Split Horizon

Routers connected to broadcast-type IP networks and using distance-vector routing protocols normally use the split-horizon mechanism to reduce the possibility of routing loops. Split horizon blocks information about routes from being advertised by a router on any interface from which that information originated. This feature usually optimizes communication among multiple routers, especially when links are broken.



Note

In general, disabling split horizon is not recommended unless you are certain that your application requires it to properly advertise routes.

If you want to configure an interface running RIP to advertise a summarized local IP address pool on a network access server for dial-up clients, use the **ip summary-address rip** interface configuration command.



Note

If split horizon is enabled, neither autosummary nor interface IP summary addresses are advertised.

Beginning in privileged EXEC mode, follow these steps to set an interface to advertise a summarized local IP address and to disable split horizon on the interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	ip address <i>ip-address subnet-mask</i>	Configure the IP address and IP subnet.
Step 4	ip summary-address rip <i>ip address ip-network mask</i>	Configure the IP address to be summarized and the IP network mask.
Step 5	no ip split horizon	Disable split horizon on the interface.
Step 6	end	Return to privileged EXEC mode.
Step 7	show ip interface <i>interface-id</i>	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable IP summarization, use the **no ip summary-address rip** router configuration command.

In this example, the major net is 10.0.0.0. The summary address 10.2.0.0 overrides the autosummary address of 10.0.0.0 so that 10.2.0.0 is advertised out Gigabit Ethernet port 2, and 10.0.0.0 is not advertised. In the example, if the interface is still in Layer 2 mode (the default), you must enter a **no switchport** interface configuration command before entering the **ip address** interface configuration command.



Note

If split horizon is enabled, neither autosummary nor interface summary addresses (those configured with the **ip summary-address rip** router configuration command) are advertised.


```
Switch(config)# router rip
Switch(config-router)# interface gigabitethernet1/0/2
Switch(config-if)# ip address 10.1.5.1 255.255.255.0
Switch(config-if)# ip summary-address rip 10.2.0.0 255.255.0.0
Switch(config-if)# no ip split-horizon
Switch(config-if)# exit
Switch(config)# router rip
Switch(config-router)# network 10.0.0.0
Switch(config-router)# neighbor 2.2.2.2 peer-group mygroup
Switch(config-router)# end
```

Configuring IGRP

Interior Gateway Routing Protocol (IGRP) is a dynamic, distance-vector routing, proprietary Cisco protocol for routing in an autonomous system (AS) that contains large, arbitrarily complex networks with diverse bandwidth and delay characteristics. IGRP uses a combination of user-configurable metrics, including internetwork delay, bandwidth, reliability, and load. IGRP also advertises types of routes: interior, system, and exterior, as shown in [Figure 28-4](#).

- *Interior routes* are routes between subnets in the network attached to a router interface. If the network attached to a router is not subnetted, IGRP does not advertise interior routes.
- *System routes* are routes to networks within an autonomous system. The router derives system routes from directly connected network interfaces and system route information provided by other IGRP-speaking routers or access servers. System routes do not include subnet information.
- *Exterior routes* are routes to networks outside the AS that are considered when identifying a gateway of last resort. The router chooses a gateway of last resort from the list of exterior routes that IGRP provides if it does not have a better route for a packet and the destination is not a connected network. If the AS has more than one connection to an external network, different routers can choose different exterior routers as the gateway of last resort.

Figure 28-4 Interior, System, and Exterior Routes



By default, a router running IGRP sends an update broadcast every 90 seconds and declares a route inaccessible if it does not receive an update from the first router in the route within three update periods (270 seconds). After seven update periods (630 seconds), the route is removed from the routing table.

This section briefly describes how to configure IGRP. It includes this information:

- [Default IGRP Configuration, page 28-24](#)
- [Understanding Load Balancing and Traffic Distribution Control, page 28-24](#)
- [Configuring Basic IGRP Parameters, page 28-25](#)
- [Configuring Split Horizon, page 28-27](#)

Default IGRP Configuration

Table 28-5 shows the default IGRP configuration.

Table 28-5 Default IGRP Configuration

Feature	Default Setting
IP split horizon	Varies with media.
Metric holddown	Disabled.
Metric maximum-hops	100 hops.
Neighbor	None defined.
Network	None specified.
Offset-list	Disabled.
Set metric	None set in route map.
Timers basic	Update: 90 seconds. Invalid: 270 seconds. Hold-down: 280 seconds. Flush: 630 seconds. Sleeptime: 0 milliseconds.
Traffic-share	Distributed proportionately to the ratios of the metrics.

Routers running IGRP use flash and poison-reverse updates to speed up the convergence of the routing algorithm. Flash updates are updates sent before the standard interval, notifying other routers of a metric change. Poison-reverse updates are intended to defeat larger routing loops caused by increases in routing metrics. The poison-reverse updates are sent to remove a route and place it in hold-down, which keeps new routing information from being used for a certain period of time.

Understanding Load Balancing and Traffic Distribution Control

IGRP can simultaneously use an asymmetric set of paths for a given destination. This unequal-cost load balancing allows traffic to be distributed among up to four unequal-cost paths to provide greater overall throughput and reliability.

Alternate path variance (that is, the difference in desirability between the primary and alternate paths) determines the feasibility of a potential route. An alternate route is feasible if the next router in the path is closer to the destination (has a lower metric value) than the router being used, and if the metric for the entire alternate path is within the variance. Only feasible paths are used for load balancing and are included in the routing table. These conditions limit the number of load balancing occurrences, but ensure that the dynamics of the network remain stable.

These general rules apply to IGRP unequal-cost load balancing:

- IGRP accepts up to four paths for a given destination network.
- The local best metric must be greater than the metric learned from the next router; that is, the next hop router must be closer (have a smaller metric value) to the destination than the local best metric.

- The alternative path metric must be within the specified variance of the local best metric. The multiplier times the local best metric for the destination must be greater than or equal to the metric through the next router.

If these conditions are met, the route is determined to be feasible and can be added to the routing table.

By default, the amount of variance is set to one (equal-cost load balancing). Use the **variance** router configuration command to define how much worse an alternate path can be before that path is disallowed.

If variance is configured as described in the preceding section, IGRP or Enhanced IGRP distributes traffic among multiple routes of unequal cost to the same destination. If you want faster convergence to alternate routes, but you do not want to send traffic across inferior routes in the normal case, you might prefer to have no traffic flow along routes with higher metrics. Use the **traffic-share** router configuration command to control distribution of traffic among multiple routes of unequal cost.

**Note**

For more information and examples, refer to the *Cisco IOS IP and IP Routing Configuration Guide for Release 12.1*.

Configuring Basic IGRP Parameters

Beginning in privileged EXEC mode, follow these steps to configure IGRP. Configuring the routing process is required; other steps are optional:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router igrp <i>autonomous-system</i>	Enable an IGRP routing process, and enter router configuration mode. The AS number identifies the routes to other IGRP routers and tags routing information.
Step 3	network <i>network-number</i>	Associate networks with an IGRP routing process. IGRP sends updates to the interfaces in the specified networks. If an interface's network is not specified, it is not advertised in any IGRP update. It is not necessary to have a registered AS number, but if you do have a registered number, we recommend that you use it to identify your process.
Step 4	offset list [<i>access-list number</i> <i>name</i>] { in out } <i>offset</i> [<i>type number</i>]	(Optional) Apply an offset list to routing metrics to increase incoming and outgoing metrics to routes learned through IGRP. You can limit the offset list with an access list or an interface.
Step 5	neighbor <i>ip-address</i>	(Optional) Define a neighboring router with which to exchange routing information. This step allows routing updates from RIP (normally a broadcast protocol) to reach nonbroadcast network.
Step 6	metric weights <i>tos k1 k2 k3 k4 k5</i>	(Optional) Adjust the IGRP metric. By default, the IGRP composite metric is a 23-bit quantity that is the sum of the segment delays and the lowest segment bandwidth for a given route. • <i>tos</i>—Type of services; the default is 0. • <i>k1-k5</i>—Constants that convert a metric vector into a scalar quantity. Defaults for k1 and k3 are 1; all others are 0.

	Command	Purpose
Step 7	timers basic <i>update invalid holddown flush [sleeptime]</i>	(Optional) Adjust routing protocol timers. <i>update</i>—The time (in seconds) between sending of routing updates. The default is 90 seconds. <i>invalid</i>—The timer interval (in seconds) after which a route is declared invalid. The default is 270 seconds. <i>holddown</i>—The time (in seconds) during which routing information about better paths is suppressed. The default is 280 seconds. <i>flush</i>—The time (in seconds) that must pass before a route is removed from the routing table. The default is 630 seconds. <i>sleeptime</i>—Interval in milliseconds for postponing routing updates. The default is 0.
Step 8	no metric holddown	(Optional) Disable the IGRP hold-down period. The route to a network is placed in holddown if the router learns that the network is farther away than previously known or is down. Holddown keeps new routing information from being used for a certain period of time. This can prevent routing loops caused by slow convergence. It is sometimes advantageous to disable holddown to increase the network's ability to quickly respond to topology changes; this command provides this function. Use the metric holddown command if other routers or access servers within the IGRP autonomous system are not configured with the no metric holddown command. If all routers are not configured the same way, you increase the possibility of routing loops.
Step 9	metric maximum-hops <i>hops</i>	(Optional) Configure the maximum network diameter. Routes with hop counts exceeding this diameter are not advertised. The default is 100 hops; the maximum is 255 hops.
Step 10	no validate-update-source	(Optional) Disable validation of the source IP address of incoming RIP routing updates. By default, the switch validates the source IP address of incoming RIP routing updates and discards the update if the source address is not valid.
Step 11	variance <i>multiplier</i>	(Optional) Define the variance associated with a particular path to enable unequal-cost load balancing if desired, balancing traffic across all feasible paths to converge to a new path if a path should fail. The multiplier can be from 1 to 128; the default is 1 (equal-cost load balancing).
Step 12	traffic-share { balanced min }	(Optional) Distribute traffic by one of these methods: balanced—Proportionately to the ratios of metrics min—By the minimum-cost route.
Step 13	end	Return to privileged EXEC mode.
Step 14	show ip protocols	Verify your entries.
Step 15	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To shut down an IGRP routing process, use the **no router igrp** global configuration command.

This example shows how to configure a router for IGRP and assign it autonomous system 109. The **network** router configuration commands show the networks directly connected to the router.

```
Switch(config)# router igrp 109
Switch(config-router)# network 131.108.0.0
Switch(config-router)# network 192.31.7.0
```

Configuring Split Horizon

Routers connected to broadcast-type IP networks and using distance-vector routing protocols normally use the split-horizon mechanism to reduce the possibility of routing loops. Split horizon blocks information about routes from being advertised by a router on any interface from which that information originated. This feature can optimize communication among multiple routers, especially when links are broken.



Note

In general, we do not recommend disabling split horizon unless you are certain that your application requires it to properly advertise routes.

Beginning in privileged EXEC mode, follow these steps to disable split horizon on the interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to configure.
Step 3	ip address <i>ip-address subnet-mask</i>	Configure the IP address and IP subnet.
Step 4	no ip split-horizon	Disable split horizon on the interface.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip interface <i>interface-id</i>	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To enable the split horizon mechanism, use the **ip split-horizon** interface configuration command.

Configuring OSPF

This section briefly describes how to configure Open Shortest Path First (OSPF). For a complete description of the OSPF commands, refer to the “OSPF Commands” chapter of the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

**Note**

OSPF classifies different media into broadcast, nonbroadcast, and point-to-point networks. The switch supports broadcast (Ethernet, Token Ring, and FDDI) and point-to-point networks (Ethernet interfaces configured as point-to-point links).

OSPF is an Interior Gateway Protocol (IGP) designed expressly for IP networks, supporting IP subnetting and tagging of externally derived routing information. OSPF also allows packet authentication and uses IP multicast when sending and receiving packets. The Cisco Systems implementation supports RFC 1253, OSPF MIB.

The Cisco implementation conforms to the OSPF Version 2 specifications with these key features:

- Definition of stub areas is supported.
- Routes learned through any IP routing protocol can be redistributed into another IP routing protocol. At the intradomain level, this means that OSPF can import routes learned through IGRP and RIP. OSPF routes can also be exported into IGRP and RIP.
- Plain text and MD5 authentication among neighboring routers within an area is supported.
- Configurable routing interface parameters include interface output cost, retransmission interval, interface transmit delay, router priority, router dead and hello intervals, and authentication key.
- Virtual links are supported.
- Not-so-stubby-areas (NSSAs) per RFC 1587 are supported.

OSPF typically requires coordination among many internal routers, *area border routers* (ABRs) connected to multiple areas, and *autonomous system boundary routers* (ASBRs). The minimum configuration would use all default parameter values, no authentication, and interfaces assigned to areas. If you customize your environment, you must ensure coordinated configuration of all routers.

This section briefly describes how to configure OSPF. It includes this information:

- [Default OSPF Configuration, page 28-29](#)
- [Configuring Basic OSPF Parameters, page 28-30](#)
- [Configuring OSPF Interfaces, page 28-31](#)
- [Configuring OSPF Area Parameters, page 28-32](#)
- [Configuring Other OSPF Parameters, page 28-33](#)
- [Changing LSA Group Pacing, page 28-35](#)
- [Configuring a Loopback Interface, page 28-35](#)
- [Monitoring OSPF, page 28-36](#)

Default OSPF Configuration

Table 28-6 shows the default OSPF configuration.

Table 28-6 Default OSPF Configuration

Feature	Default Setting
Interface parameters	Cost: No default cost predefined. Retransmit interval: 5 seconds. Transmit delay: 1 second. Priority: 1. Hello interval: 10 seconds. Dead interval: 4 times the hello interval. No authentication. No password specified. MD5 authentication disabled.
Area	Authentication type: 0 (no authentication). Default cost: 1. Range: Disabled. Stub: No stub area defined. NSSA: No NSSA area defined.
Auto cost	100 Mbps.
Default-information originate	Disabled. When enabled, the default metric setting is 10, and the external route type default is Type 2.
Default metric	Built-in, automatic metric translation, as appropriate for each routing protocol.
Distance OSPF	dist1 (all routes within an area): 110. dist2 (all routes from one area to another): 110. and dist3 (routes from other routing domains): 110.
OSPF database filter	Disabled. All outgoing link-state advertisements (LSAs) are flooded to the interface.
IP OSPF name lookup	Disabled.
Log adjacency changes	Enabled.
Neighbor	None specified.
Neighbor database filter	Disabled. All outgoing LSAs are flooded to the neighbor.
Network area	Disabled.
Router ID	No OSPF routing process defined.
Summary address	Disabled.
Timers LSA group pacing	240 seconds.

Table 28-6 *Default OSPF Configuration (continued)*

Feature	Default Setting
Timers shortest path first (spf)	spf delay: 5 seconds. spf-holdtime: 10 seconds.
Virtual link	No area ID or router ID defined. Hello interval: 10 seconds. Retransmit interval: 5 seconds. Transmit delay: 1 second. Dead interval: 40 seconds. Authentication key: no key predefined. Message-digest key (MD5): no key predefined.

Configuring Basic OSPF Parameters

Enabling OSPF requires that you create an OSPF routing process, specify the range of IP addresses to be associated with the routing process, and assign area IDs to be associated with that range.

Beginning in privileged EXEC mode, follow these steps to enable OSPF:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router ospf <i>process-id</i>	Enable OSPF routing, and enter router configuration mode. The process ID is an internally used identification parameter that is locally assigned and can be any positive integer. Each OSPF routing process has a unique value.
Step 3	network <i>address wildcard-mask area</i> <i>area-id</i>	Define an interface on which OSPF runs and the area ID for that interface. You can use the wildcard-mask to use a single command to define one or more multiple interfaces to be associated with a specific OSPF area. The area ID can be a decimal value or an IP address.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip protocols	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To terminate an OSPF routing process, use the **no router ospf** *process-id* global configuration command.

This example shows how to configure an OSPF routing process and assign it a process number of 109:

```
Switch(config)# router ospf 109
Switch(config-router)# network 131.108.0.0 255.255.255.0 area 24
```

Configuring OSPF Interfaces

You can use the **ip ospf** interface configuration commands to modify interface-specific OSPF parameters. You are not required to modify any of these parameters, but some interface parameters (hello interval, dead interval, and authentication key) must be consistent across all routers in an attached network. If you modify these parameters, be sure all routers in the network have compatible values.



Note

The **ip ospf** interface configuration commands are all optional.

Beginning in privileged EXEC mode, follow these steps to modify OSPF interface parameters:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	ip ospf cost	(Optional) Explicitly specify the cost of sending a packet on the interface.
Step 4	ip ospf retransmit-interval <i>seconds</i>	(Optional) Specify the number of seconds between link state advertisement transmissions. The range is 1 to 65535 seconds. The default is 5 seconds.
Step 5	ip ospf transmit-delay <i>seconds</i>	(Optional) Set the estimated number of seconds to wait before sending a link state update packet. The range is 1 to 65535 seconds. The default is 1 second.
Step 6	ip ospf priority <i>number</i>	(Optional) Set priority to help determine the OSPF designated router for a network. The range is from 0 to 255. The default is 1.
Step 7	ip ospf hello-interval <i>seconds</i>	(Optional) Set the number of seconds between hello packets sent on an OSPF interface. The value must be the same for all nodes on a network. The range is 1 to 65535 seconds. The default is 10 seconds.
Step 8	ip ospf dead-interval <i>seconds</i>	(Optional) Set the number of seconds after the last device hello packet was seen before its neighbors declare the OSPF router to be down. The value must be the same for all nodes on a network. The range is 1 to 65535 seconds. The default is 4 times the hello interval.
Step 9	ip ospf authentication-key <i>key</i>	(Optional) Assign a password to be used by neighboring OSPF routers. The password can be any string of keyboard-entered characters up to 8 bytes in length. All neighboring routers on the same network must have the same password to exchange OSPF information.
Step 10	ip ospf message digest-key <i>keyid md5 key</i>	(Optional) Enable MDS authentication. <i>keyid</i>—An identifier from 1 to 255. <i>key</i>—An alphanumeric password of up to 16 bytes.
Step 11	ip ospf database-filter all out	(Optional) Block flooding of OSPF LSA packets to the interface. By default, OSPF floods new LSAs over all interfaces in the same area, except the interface on which the LSA arrives.
Step 12	end	Return to privileged EXEC mode.

	Command	Purpose
Step 13	show ip ospf interface [<i>interface-name</i>]	Display OSPF-related interface information.
Step 14	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no** form of these commands to remove the configured parameter value or return to the default value.

Configuring OSPF Area Parameters

You can optionally configure several OSPF area parameters. These parameters include authentication for password-based protection against unauthorized access to an area, stub areas, and not-so-stubby-areas (NSSAs). *Stub areas* are areas into which information on external routes is not sent. Instead, the area border router (ABR) generates a default external route into the stub area for destinations outside the autonomous system (AS). An NSSA does not flood all LSAs from the core into the area, but can import AS external routes within the area by redistribution.

Route summarization is the consolidation of advertised addresses into a single summary route to be advertised by other areas. If network numbers are contiguous, you can use the **area range** router configuration command to configure the ABR to advertise a summary route that covers all networks in the range.



Note

The OSPF **area** router configuration commands are all optional.

Beginning in privileged EXEC mode, follow these steps to configure area parameters:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router ospf <i>process-id</i>	Enable OSPF routing, and enter router configuration mode.
Step 3	area <i>area-id</i> authentication	(Optional) Allow password-based protection against unauthorized access to the identified area. The identifier can be either a decimal value or an IP address.
Step 4	area <i>area-id</i> authentication message-digest	(Optional) Enable MD5 authentication on the area.
Step 5	area <i>area-id</i> stub [no-summary]	(Optional) Define an area as a stub area. The no-summary keyword prevents an ABR from sending summary link advertisements into the stub area.
Step 6	area <i>area-id</i> nssa [no-redistribution] [default-information-originate] [no-summary]	(Optional) Defines an area as a not-so-stubby-area. Every router within the same area must agree that the area is NSSA. Select one of these keywords: no-redistribution—Select when the router is an NSSA ABR and you want the redistribute command to import routes into normal areas, but not into the NSSA. default-information-originate—Select on an ABR to allow importing type 7 LSAs into the NSSA. no-redistribution—Select to not send summary LSAs into the NSSA.

	Command	Purpose
Step 7	area <i>area-id</i> range <i>address mask</i>	(Optional) Specify an address range for which a single route is advertised. Use this command only with area border routers.
Step 8	end	Return to privileged EXEC mode.
Step 9	show ip ospf [<i>process-id</i>]	Display information about the OSPF routing process in general or for a specific process ID to verify configuration.
	show ip ospf [<i>process-id</i> [<i>area-id</i>]] database	Display lists of information related to the OSPF database for a specific router.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no** form of these commands to remove the configured parameter value or to return to the default value.

Configuring Other OSPF Parameters

You can optionally configure other OSPF parameters in router configuration mode.

- **Route summarization:** When redistributing routes from other protocols as described in the [“Using Route Maps to Redistribute Routing Information”](#) section on page 28-90, each route is advertised individually in an external LSA. To help decrease the size of the OSPF link state database, you can use the **summary-address** router configuration command to advertise a single router for all the redistributed routes included in a specified network address and mask.
- **Virtual links:** In OSPF, all areas must be connected to a backbone area. You can establish a virtual link in case of a backbone-continuity break by configuring two Area Border Routers as endpoints of a virtual link. Configuration information includes the identity of the other virtual endpoint (the other ABR) and the nonbackbone link that the two routers have in common (the transit area). Virtual links cannot be configured through a stub area.
- **Default route:** When you specifically configure redistribution of routes into an OSPF routing domain, the route automatically becomes an autonomous system boundary router (ASBR). You can force the ASBR to generate a default route into the OSPF routing domain.
- **Domain Name Server (DNS) names** for use in all OSPF **show** privileged EXEC command displays makes it easier to identify a router than displaying it by router ID or neighbor ID.
- **Default Metrics:** OSPF calculates the OSPF metric for an interface according to the bandwidth of the interface. The metric is calculated as *ref-bw* divided by bandwidth, where *ref* is 10 by default, and bandwidth (*bw*) is determined by the **bandwidth** interface configuration command. For multiple links with high bandwidth, you can specify a larger number to differentiate the cost on those links.
- **Administrative distance** is a rating of the trustworthiness of a routing information source, an integer between 0 and 255, with a higher value meaning a lower trust rating. An administrative distance of 255 means the routing information source cannot be trusted at all and should be ignored. OSPF uses three different administrative distances: routes within an area (interarea), routes to another area (interarea), and routes from another routing domain learned through redistribution (external). You can change any of the distance values.
- **Passive interfaces:** Because interfaces between two devices on an Ethernet represent only one network segment, to prevent OSPF from sending hello packets for the sending interface, you must configure the sending device to be a passive interface. Both devices can identify each other through the hello packet for the receiving interface.

- Route calculation timers: You can configure the delay time between when OSPF receives a topology change and when it starts the shortest path first (SPF) calculation and the hold time between two SPF calculations.
- Log neighbor changes: You can configure the router to send a syslog message when an OSPF neighbor state changes, providing a high-level view of changes in the router.

Beginning in privileged EXEC mode, follow these steps to configure these OSPF parameters:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router ospf <i>process-id</i>	Enable OSPF routing, and enter router configuration mode.
Step 3	summary-address <i>address mask</i>	(Optional) Specify an address and IP subnet mask for redistributed routes so that only one summary route is advertised.
Step 4	area <i>area-id</i> virtual-link router-id [hello-interval <i>seconds</i>] [retransmit-interval <i>seconds</i>] [trans] [[authentication-key <i>key</i>] message-digest-key <i>keyid md5 key</i>]]	(Optional) Establish a virtual link and set its parameters. See the “Configuring OSPF Interfaces” section on page 28-31 for parameter definitions and Table 28-6 on page 28-29 for virtual link defaults.
Step 5	default-information originate [always] [metric <i>metric-value</i>] [metric-type <i>type-value</i>] [route-map <i>map-name</i>]	(Optional) Force the ASBR to generate a default route into the OSPF routing domain. Parameters are all optional.
Step 6	ip ospf name-lookup	(Optional) Configure DNS name lookup. The default is disabled.
Step 7	ip auto-cost reference-bandwidth <i>ref-bw</i>	(Optional) Specify an address range for which a single route will be advertised. Use this command only with area border routers.
Step 8	distance ospf {[inter-area <i>dist1</i>] [inter-area <i>dist2</i>] [external <i>dist3</i>]}	(Optional) Change the OSPF distance values. The default distance for each type of route is 110. The range is 1 to 255.
Step 9	passive-interface <i>type number</i>	(Optional) Suppress the sending of hello packets through the specified interface.
Step 10	timers spf <i>spf-delay spf-holdtime</i>	(Optional) Configure route calculation timers. • <i>spf-delay</i>—Enter an integer from 0 to 65535. The default is 5 seconds; 0 means no delay. • <i>spf-holdtime</i>—Enter an integer from 0 to 65535. The default is 10 seconds; 0 means no delay.
Step 11	ospf log-adj-changes	(Optional) Send syslog message when a neighbor state changes.
Step 12	end	Return to privileged EXEC mode.
Step 13	show ip ospf [<i>process-id</i> [<i>area-id</i>]] database	Display lists of information related to the OSPF database for a specific router. For some of the keyword options, see to the “Monitoring OSPF” section on page 28-36 .
Step 14	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Changing LSA Group Pacing

The OSPF LSA group pacing feature allows the router to group OSPF LSAs and pace the refreshing, check-summing, and aging functions for more efficient router use. This feature is enabled by default with a 4-minute default pacing interval, and you will not usually need to modify this parameter. The optimum group pacing interval is inversely proportional to the number of LSAs the router is refreshing, check-summing, and aging. For example, if you have approximately 10,000 LSAs in the database, decreasing the pacing interval would benefit you. If you have a very small database (40 to 100 LSAs), increasing the pacing interval to 10 to 20 minutes might benefit you slightly.

Beginning in privileged EXEC mode, follow these steps to configure OSPF LSA pacing:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router ospf <i>process-id</i>	Enable OSPF routing, and enter router configuration mode.
Step 3	timers lsa-group-pacing <i>seconds</i>	Change the group pacing of LSAs.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default value, use the **no timers lsa-group-pacing** router configuration command.

Configuring a Loopback Interface

OSPF uses the highest IP address configured on the interfaces as its router ID. If this interface is down or removed, the OSPF process must recalculate a new router ID and resend all its routing information out its interfaces. If a loopback interface is configured with an IP address, OSPF uses this IP address as its router ID, even if other interfaces have higher IP addresses. Because loopback interfaces never fail, this provides greater stability. OSPF automatically prefers a loopback interface over other interfaces, and it chooses the highest IP address among all loopback interfaces.

Beginning in privileged EXEC mode, follow these steps to configure a loopback interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface loopback 0	Create a loopback interface, and enter interface configuration mode.
Step 3	ip address <i>address mask</i>	Assign an IP address to this interface.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip interface	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no interface loopback 0** global configuration command to disable the loopback interface.

Monitoring OSPF

You can display specific statistics such as the contents of IP routing tables, caches, and databases.

Table 28-7 lists some of the privileged EXEC commands for displaying statistics. For more **show ip ospf database** privileged EXEC command options and for explanations of fields in the resulting display, refer to the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

Table 28-7 Show IP OSPF Statistics Commands

Command	Purpose
show ip ospf [<i>process-id</i>]	Display general information about OSPF routing processes.
show ip ospf [<i>process-id</i>] database [router] [<i>link-state-id</i>] show ip ospf [<i>process-id</i>] database [router] [self-originate] show ip ospf [<i>process-id</i>] database [router] [adv-router [<i>ip-address</i>]] show ip ospf [<i>process-id</i>] database [network] [<i>link-state-id</i>] show ip ospf [<i>process-id</i>] database [summary] [<i>link-state-id</i>] show ip ospf [<i>process-id</i>] database [asbr-summary] [<i>link-state-id</i>] show ip ospf [<i>process-id</i>] database [external] [<i>link-state-id</i>] show ip ospf [<i>process-id area-id</i>] database [database-summary]	Display lists of information related to the OSPF database.
show ip ospf border-routes	Display the internal OSPF routing ABR and ASBR table entries.
show ip ospf interface [<i>interface-name</i>]	Display OSPF-related interface information.
show ip ospf neighbor [<i>interface-name</i>] [<i>neighbor-id</i>] detail	Display OSPF interface neighbor information.
show ip ospf virtual-links	Display OSPF-related virtual links information.

Configuring EIGRP

Enhanced IGRP (EIGRP) is a Cisco proprietary enhanced version of the IGRP. EIGRP uses the same distance vector algorithm and distance information as IGRP; however, the convergence properties and the operating efficiency of EIGRP are significantly improved.

The convergence technology employs an algorithm referred to as the Diffusing Update Algorithm (DUAL), which guarantees loop-free operation at every instant throughout a route computation and allows all devices involved in a topology change to synchronize at the same time. Routers that are not affected by topology changes are not involved in recomputations.

IP EIGRP provides increased network width. With RIP, the largest possible width of your network is 15 hops. When IGRP is enabled, the largest possible width is 224 hops. Because the EIGRP metric is large enough to support thousands of hops, the only barrier to expanding the network is the transport-layer hop counter. EIGRP increments the transport control field only when an IP packet has traversed 15 routers and the next hop to the destination was learned through EIGRP. When a RIP route is used as the next hop to the destination, the transport control field is incremented as usual.

EIGRP offers these features:

- Fast convergence.
- Incremental updates when the state of a destination changes, instead of sending the entire contents of the routing table, minimizing the bandwidth required for EIGRP packets.
- Less CPU usage than IGRP because full update packets need not be processed each time they are received.
- Protocol-independent neighbor discovery mechanism to learn about neighboring routers.
- Variable-length subnet masks (VLSMs).
- Arbitrary route summarization.
- EIGRP scales to large networks.

Enhanced IGRP has these four basic components:

- *Neighbor discovery and recovery* is the process that routers use to dynamically learn of other routers on their directly attached networks. Routers must also discover when their neighbors become unreachable or inoperative. Neighbor discovery and recovery is achieved with low overhead by periodically sending small hello packets. As long as hello packets are received, the Cisco IOS software can determine that a neighbor is alive and functioning. When this status is determined, the neighboring routers can exchange routing information.
- *The reliable transport protocol* is responsible for guaranteed, ordered delivery of EIGRP packets to all neighbors. It supports intermixed transmission of multicast and unicast packets. Some EIGRP packets must be sent reliably, and others need not be. For efficiency, reliability is provided only when necessary. For example, on a multiaccess network that has multicast capabilities (such as Ethernet), it is not necessary to send hellos reliably to all neighbors individually. Therefore, EIGRP sends a single multicast hello with an indication in the packet informing the receivers that the packet need not be acknowledged. Other types of packets (such as updates) require acknowledgment, which is shown in the packet. The reliable transport has a provision to send multicast packets quickly when there are unacknowledged packets pending. Doing so helps ensure that convergence time remains low in the presence of varying speed links.
- *The DUAL finite state machine* embodies the decision process for all route computations. It tracks all routes advertised by all neighbors. DUAL uses the distance information (known as a metric) to select efficient, loop-free paths. DUAL selects routes to be inserted into a routing table based on feasible successors. A successor is a neighboring router used for packet forwarding that has a least-cost path to a destination that is guaranteed not to be part of a routing loop. When there are no

feasible successors, but there are neighbors advertising the destination, a recomputation must occur. This is the process whereby a new successor is determined. The amount of time it takes to recompute the route affects the convergence time. Recomputation is processor-intensive; it is advantageous to avoid recomputation if it is not necessary. When a topology change occurs, DUAL tests for feasible successors. If there are feasible successors, it uses any it finds to avoid unnecessary recomputation.

- The *protocol-dependent modules* are responsible for network layer protocol-specific tasks. An example is the IP EIGRP module, which is responsible for sending and receiving EIGRP packets that are encapsulated in IP. It is also responsible for parsing EIGRP packets and informing DUAL of the new information received. EIGRP asks DUAL to make routing decisions, but the results are stored in the IP routing table. EIGRP is also responsible for redistributing routes learned by other IP routing protocols.

This section briefly describes how to configure EIGRP. It includes this information:

- [Default EIGRP Configuration, page 28-38](#)
- [Configuring Basic EIGRP Parameters, page 28-39](#)
- [Configuring EIGRP Interfaces, page 28-40](#)
- [Configuring EIGRP Route Authentication, page 28-41](#)
- [Monitoring and Maintaining EIGRP, page 28-42](#)

Default EIGRP Configuration

[Table 28-8](#) shows the default EIGRP configuration.

Table 28-8 Default EIGRP Configuration

Feature	Default Setting
Auto summary	Enabled. Subprefixes are summarized to the classful network boundary when crossing classful network boundaries.
Default-information	Exterior routes are accepted and default information is passed between IGRP or EIGRP processes when doing redistribution.
Default metric	Only connected routes and interface static routes can be redistributed without a default metric. The metric includes: • Bandwidth: 0 or greater kbps. • Delay (tens of microseconds): 0 or any positive number that is a multiple of 39.1 nanoseconds. • Reliability: any number between 0 and 255 (255 means 100 percent reliability). • Loading: effective bandwidth as a number between 0 and 255 (255 is 100 percent loading). • MTU: maximum transmission unit size of the route in bytes. 0 or any positive integer.
Distance	Internal distance: 90. External distance: 170.
EIGRP log-neighbor changes	Disabled. No adjacency changes logged.
IP authentication key-chain	No authentication provided.

Table 28-8 Default EIGRP Configuration (continued)

Feature	Default Setting
IP authentication mode	No authentication provided.
IP bandwidth-percent	50 percent.
IP hello interval	For low-speed nonbroadcast multiaccess (NBMA) networks: 60 seconds; all other networks: 5 seconds.
IP hold-time	For low-speed NBMA networks: 180 seconds; all other networks: 15 seconds.
IP split-horizon	Enabled.
IP summary address	No summary aggregate addresses are predefined.
Metric weights	tos: 0; k1 and k3: 1; k2, k4, and k5: 0
Network	None specified.
Offset-list	Disabled.
Router EIGRP	Disabled.
Set metric	No metric set in the route map.
Traffic-share	Distributed proportionately to the ratios of the metrics.
Variance	1 (equal-cost load balancing).

To create an EIGRP routing process, you must enable EIGRP and associate networks. EIGRP sends updates to the interfaces in the specified networks. If you do not specify an interface network, it is not advertised in any EIGRP update.

**Note**

If you have routers on your network that are configured for IGRP, and you want to change to EIGRP, you must designate transition routers that have both IGRP and EIGRP configured. In these cases, perform Steps 1 through 3 in the next section and also see the [“Configuring IGRP” section on page 28-23](#). You must use the same AS number for routes to be automatically redistributed.

Configuring Basic EIGRP Parameters

Beginning in privileged EXEC mode, follow these steps to configure EIGRP. Configuring the routing process is required; other steps are optional:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router eigrp <i>autonomous-system</i>	Enable an EIGRP routing process, and enter router configuration mode. The AS number identifies the routes to other EIGRP routers and is used to tag routing information.
Step 3	network <i>network-number</i>	Associate networks with an EIGRP routing process. EIGRP sends updates to the interfaces in the specified networks. If an interface's network is not specified, it is not advertised in any IGRP or EIGRP update.

	Command	Purpose
Step 4	eigrp log-neighbor-changes	(Optional) Enable logging of EIGRP neighbor changes to monitor routing system stability.
Step 5	metric weights <i>tos k1 k2 k3 k4 k5</i>	(Optional) Adjust the EIGRP metric. Although the defaults have been carefully determined to provide excellent operation in most networks, you can adjust them.
		
	Caution	Determining metrics is complex and is not recommended without guidance from an experienced network designer.
Step 6	offset list [<i>access-list number</i> <i>name</i>] { in out } <i>offset</i> [<i>type number</i>]	(Optional) Apply an offset list to routing metrics to increase incoming and outgoing metrics to routes learned through EIGRP. You can limit the offset list with an access list or an interface.
Step 7	no auto-summary	(Optional) Disable automatic summarization of subnet routes into network-level routes.
Step 8	ip summary-address eigrp <i>autonomous-system-number address mask</i>	(Optional) Configure a summary aggregate.
Step 9	end	Return to privileged EXEC mode.
Step 10	show ip protocols	Verify your entries.
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no** forms of these commands to disable the feature or return the setting to the default value.

Configuring EIGRP Interfaces

Other optional EIGRP parameters can be configured on an interface basis.

Beginning in privileged EXEC mode, follow these steps to configure EIGRP interfaces:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	ip bandwidth-percent eigrp <i>percent</i>	(Optional) Configure the percentage of bandwidth that can be used by EIGRP on an interface. The default is 50 percent.
Step 4	ip summary-address eigrp <i>autonomous-system-number address mask</i>	(Optional) Configure a summary aggregate address for a specified interface (not usually necessary if auto-summary is enabled).
Step 5	ip hello-interval eigrp <i>autonomous-system-number seconds</i>	(Optional) Change the hello time interval for an EIGRP routing process. The range is 1 to 65535 seconds. The default is 60 seconds for low-speed NBMA networks and 5 seconds for all other networks.

	Command	Purpose
Step 6	ip hold-time eigrp <i>autonomous-system-number</i> <i>seconds</i>	(Optional) Change the hold time interval for an EIGRP routing process. The range is 1 to 65535 seconds. The default is 180 seconds for low-speed NBMA networks and 15 seconds for all other networks.
		 Caution Do not adjust the hold time without consulting Cisco technical support.
Step 7	no ip split-horizon eigrp <i>autonomous-system-number</i>	(Optional) Disable split horizon to allow route information to be advertised by a router out any interface from which that information originated.
Step 8	end	Return to privileged EXEC mode.
Step 9	show ip eigrp interface	Display which interfaces EIGRP is active on and information about EIGRP relating to those interfaces.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no** forms of these commands to disable the feature or return the setting to the default value.

Configuring EIGRP Route Authentication

EIGRP route authentication provides MD5 authentication of routing updates from the EIGRP routing protocol to prevent the introduction of unauthorized or false routing messages from unapproved sources.

Beginning in privileged EXEC mode, follow these steps to enable authentication:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	ip authentication mode eigrp <i>autonomous-system</i> md5	Enable MD5 authentication in IP EIGRP packets.
Step 4	ip authentication key-chain eigrp <i>autonomous-system</i> <i>key-chain</i>	Enable authentication of IP EIGRP packets.
Step 5	exit	Return to global configuration mode.
Step 6	key chain <i>name-of-chain</i>	Identify a key chain and enter key-chain configuration mode. Match the name configured in Step 4.
Step 7	key <i>number</i>	In key-chain configuration mode, identify the key number.
Step 8	key-string <i>text</i>	In key-chain key configuration mode, identify the key string.

	Command	Purpose
Step 9	accept-lifetime <i>start-time</i> { infinite <i>end-time</i> duration <i>seconds</i> }	(Optional) Specify the time period during which the key can be received. The <i>start-time</i> and <i>end-time</i> syntax can be either <i>hh:mm:ss Month date year</i> or <i>hh:mm:ss date Month year</i> . The default is forever with the default <i>start-time</i> and the earliest acceptable date as January 1, 1993. The default <i>end-time</i> and duration is infinite .
Step 10	send-lifetime <i>start-time</i> { infinite <i>end-time</i> duration <i>seconds</i> }	(Optional) Specify the time period during which the key can be sent. The <i>start-time</i> and <i>end-time</i> syntax can be either <i>hh:mm:ss Month date year</i> or <i>hh:mm:ss date Month year</i> . The default is forever with the default <i>start-time</i> and the earliest acceptable date as January 1, 1993. The default <i>end-time</i> and duration is infinite .
Step 11	end	Return to privileged EXEC mode.
Step 12	show key chain	Display authentication key information.
Step 13	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no** forms of these commands to disable the feature or to return the setting to the default value.

Monitoring and Maintaining EIGRP

You can delete neighbors from the neighbor table. You can also display various EIGRP routing statistics. [Table 28-9](#) lists the privileged EXEC commands for deleting neighbors and displaying statistics. For explanations of fields in the resulting display, refer to the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

Table 28-9 IP EIGRP Clear and Show Commands

Command	Purpose
clear ip eigrp neighbors [<i>if-address</i> <i>interface</i>]	Delete neighbors from the neighbor table.
show ip eigrp interface [<i>interface</i>] [<i>as number</i>]	Display information about interfaces configured for EIGRP.
show ip eigrp neighbors [<i>type-number</i>]	Display EIGRP discovered neighbors.
show ip eigrp topology [<i>autonomous-system-number</i>] [<i>ip-address</i>] <i>mask</i>]	Display the EIGRP topology table for a given process.
show ip eigrp traffic [<i>autonomous-system-number</i>]	Display the number of packets sent and received for all or a specified EIGRP process.

Configuring BGP

The Border Gateway Protocol (BGP) is an exterior gateway protocol used to set up an interdomain routing system that guarantees the loop-free exchange of routing information between autonomous systems. Autonomous systems are made up of routers that operate under the same administration and that run Interior Gateway Protocols (IGPs), such as RIP or OSPF, within their boundaries and that interconnect by using an Exterior Gateway Protocol (EGP). BGP version 4 is the standard EGP for interdomain routing in the Internet. The protocol is defined in RFCs 1163, 1267, and 1771. You can find detailed information about BGP in *Internet Routing Architectures*, published by Cisco Press, and in the “Configuring BGP” chapter in the *Cisco IOS IP and IP Routing Configuration Guide*.




Note

For details about BGP commands and keywords, refer to the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*. For a list of BGP commands that are visible but not supported by the switch, see [Appendix C, “Unsupported Commands in Cisco IOS Release 12.1\(14\)AX.”](#)

Routers that belong to the same autonomous system (AS) and that exchange BGP updates run *internal BGP* (IBGP), and routers that belong to different autonomous systems and that exchange BGP updates run *external BGP* (EBGP). Most configuration commands are the same for configuring EBGP and IBGP. The difference is that the routing updates are exchanged either between autonomous systems (EBGP) or within an AS (IBGP). [Figure 28-5](#) shows a network that is running both EBGP and IBGP.

Figure 28-5 EBGP, IBGP, and Multiple Autonomous Systems



Before exchanging information with an external AS, BGP ensures that networks within the AS can be reached by defining internal BGP peering among routers within the AS and by redistributing BGP routing information to IGPs that run within the AS, such as IGRP and OSPF.

Routers that run a BGP routing process are often referred to as BGP *speakers*. BGP uses the Transmission Control Protocol (TCP) as its transport protocol (specifically port 179). Two BGP speakers that have a TCP connection to each other for exchanging routing information are known as *peers* or *neighbors*. In [Figure 28-5](#), Routers A and B are BGP peers, as are Routers B and C and Routers C and D. The routing information is a series of AS numbers that describe the full path to the destination network. BGP uses this information to construct a loop-free map of autonomous systems.

The network has these characteristics:

- Routers A and B are running EBGp, and Routers B and C are running IBGP. Note that the EBGp peers are directly connected and that the IBGP peers are not. As long as there is an IGP running that allows the two neighbors to reach one another, IBGP peers do not have to be directly connected.
- All BGP speakers within an AS must establish a peer relationship with each other. That is, the BGP speakers within an AS must be fully meshed logically. BGP4 provides two techniques that reduce the requirement for a logical full mesh: *confederations* and *route reflectors*.
- AS 200 is a *transit AS* for AS 100 and AS 300—that is, AS 200 is used to transfer packets between AS 100 and AS 300.

BGP peers initially exchange their full BGP routing tables and then send only incremental updates. BGP peers also exchange keepalive messages (to ensure that the connection is up) and notification messages (in response to errors or special conditions).

In BGP, each route consists of a network number, a list of autonomous systems that information has passed through (the *autonomous system path*), and a list of other *path attributes*. The primary function of a BGP system is to exchange network reachability information, including information about the list of AS paths, with other BGP systems. This information can be used to determine AS connectivity, to prune routing loops, and to enforce AS-level policy decisions.

A router or switch running Cisco IOS does not select or use an IBGP route unless it has a route available to the next-hop router and it has received synchronization from an IGP (unless IGP synchronization is disabled). When multiple routes are available, BGP bases its path selection on *attribute* values. See the [“Configuring BGP Decision Attributes” section on page 28-50](#) for information about BGP attributes.

BGP Version 4 supports classless interdomain routing (CIDR) so you can reduce the size of your routing tables by creating aggregate routes, resulting in *supernets*. CIDR eliminates the concept of network classes within BGP and supports the advertising of IP prefixes.

These sections briefly describe how to configure BGP and supported BGP features:

- [Default BGP Configuration, page 28-45](#)
- [Enabling BGP Routing, page 28-47](#)
- [Managing Routing Policy Changes, page 28-49](#)
- [Configuring BGP Decision Attributes, page 28-50](#)
- [Configuring BGP Filtering with Route Maps, page 28-52](#)
- [Configuring BGP Filtering by Neighbor, page 28-53](#)
- [Configuring Prefix Lists for BGP Filtering, page 28-54](#)
- [Configuring BGP Community Filtering, page 28-55](#)
- [Configuring BGP Neighbors and Peer Groups, page 28-56](#)
- [Configuring Aggregate Addresses, page 28-58](#)
- [Configuring Routing Domain Confederations, page 28-59](#)
- [Configuring BGP Route Reflectors, page 28-60](#)
- [Configuring Route Dampening, page 28-61](#)
- [Monitoring and Maintaining BGP, page 28-62](#)

For detailed descriptions of BGP configuration, refer to the “Configuring BGP” chapter in the *Cisco IOS IP and IP Routing Configuration Guide*. For details about specific commands, refer to the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

For a list of BGP commands that are visible but not supported by the switch, see [Appendix C](#), “Unsupported Commands in Cisco IOS Release 12.1(14)AX.”

Default BGP Configuration

[Table 28-10](#) shows the basic default BGP configuration. For the defaults for all characteristics, refer to the specific commands in the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

Table 28-10 Default BGP Configuration

Feature	Default Setting
Aggregate address	Disabled: None defined.
AS path access list	None defined.
Auto summary	Enabled.
Best path	The router considers <i>as-path</i> in choosing a route and does not compare similar routes from external BGP peers. Compare router ID: Disabled.
BGP community list	Number: None defined. When you permit a value for the community number, the list defaults to an implicit deny for everything else that has not been permitted. Format: Cisco default format (32-bit number).
BGP confederation identifier/peers	Identifier: None configured. Peers: None identified.
BGP Fast external fallover	Enabled.
BGP local preference	100. The range is 0 to 4294967295 with the higher value preferred.
BGP network	None specified; no backdoor route advertised.
BGP route dampening	Disabled by default. When enabled: - Half-life is 15 minutes. - Re-use is 750 (10-second increments). - Suppress is 2000 (10-second increments). - Max-suppress-time is 4 times half-life; 60 minutes.
BGP router ID	The IP address of a loopback interface if one is configured or the highest IP address configured for a physical interface on the router.
Default information originate (protocol or network redistribution)	Disabled.
Default metric	Built-in, automatic metric translations.
Distance	External route administrative distance: 20 (acceptable values are from 1 to 255). Internal route administrative distance: 200 (acceptable values are from 1 to 255). Local route administrative distance: 200 (acceptable values are from 1 to 255).
Distribute list	In (filter networks received in updates): Disabled. Out (suppress networks from being advertised in updates): Disabled.
Internal route redistribution	Disabled.

Table 28-10 Default BGP Configuration (continued)

Feature	Default Setting
IP prefix list	None defined.
Multi exit discriminator (MED)	Always compare: Disabled. Does not compare MEDs for paths from neighbors in different autonomous systems. Best path compare: Disabled. MED missing as worst path: Disabled. Deterministic MED comparison is disabled.
Neighbor	These are default neighbor settings: • Advertisement interval: 30 seconds for external peers; 5 seconds for internal peers. • Change logging: Enabled. • Conditional advertisement: Disabled. • Default originate: No default route is sent to the neighbor. • Description: None. • Distribute list: None defined. • External BGP multihop: Only directly connected neighbors are allowed. • Filter list: None used. • Maximum number of prefixes received: No limit. • Next hop (router as next hop for BGP neighbor): Disabled. • Password: Disabled. • Peer group: None defined; no members assigned. • Prefix list: None specified. • Remote AS (add entry to neighbor BGP table): No peers defined. • Private AS number removal: Disabled. • Route maps: None applied to a peer. • Send community attributes: None sent to neighbors. • Shutdown or soft reconfiguration: Not enabled. • Timers: keepalive: 60 seconds; holdtime: 180 seconds. • Update source: Best local address. • Version: BGP version 4. • Weight: Routes learned through BGP peer: 0; routes sourced by the local router: 32768.
Route reflector	None configured.
Synchronization (BGP and IGP)	Enabled.
Table map update	Disabled.
Timers	Keepalive: 60 seconds Holdtime: 180 seconds.

Enabling BGP Routing

To enable BGP routing, you establish a BGP routing process and define the local network. Because BGP must completely recognize the relationships with its neighbors, you must also specify a BGP neighbor.

BGP supports two kinds of neighbors: internal and external. *Internal neighbors* are in the same AS; *external neighbors* are in different autonomous systems. External neighbors are usually adjacent to each other and share a subnet, but internal neighbors can be anywhere in the same AS.

The switch supports the use of private AS numbers, usually assigned by service providers and given to systems whose routes are not advertised to external neighbors. The private AS numbers are from 64512 to 65535. You can configure external neighbors to remove private AS numbers from the AS path by using the **neighbor remove-private-as** router configuration command. Then when an update is passed to an external neighbor, if the AS path includes private AS numbers, these numbers are dropped.

If your AS will be passing traffic through it from another AS to a third AS, it is important to be consistent about the routes it advertises. If BGP advertised a route before all routers in the network had learned about the route through the IGP, the AS might receive traffic that some routers could not yet route. To prevent this from happening, BGP must wait until the IGP has propagated information across the AS so that BGP is *synchronized* with the IGP. Synchronization is enabled by default. If your AS does not pass traffic from one AS to another AS, or if all routers in your autonomous systems are running BGP, you can disable synchronization, which allows your network to carry fewer routes in the IGP and allows BGP to converge more quickly.

Beginning in privileged EXEC mode, follow these steps to enable BGP routing, establish a BGP routing process, and specify a neighbor:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip routing	Enable IP routing (required only if IP routing is disabled).
Step 3	router bgp <i>autonomous-system</i>	Enable a BGP routing process, assign it an AS number, and enter router configuration mode. The AS number can be from 1 to 65535, with 64512 to 65535 designated as private autonomous numbers.
Step 4	network <i>network-number</i> [mask <i>network-mask</i>] [route-map <i>route-map-name</i>]	Configure a network as local to this AS, and enter it in the BGP table.
Step 5	neighbor { <i>ip-address</i> <i>peer-group-name</i> } remote-as <i>number</i>	Add an entry to the BGP neighbor table specifying that the neighbor identified by the IP address belongs to the specified AS. For EBGP, neighbors are usually directly connected, and the IP address is the address of the interface at the other end of the connection. For IBGP, the IP address can be the address of any of the router interfaces.
Step 6	neighbor { <i>ip-address</i> <i>peer-group-name</i> } remove-private-as	(Optional) Remove private AS numbers from the AS-path in outbound routing updates.
Step 7	no synchronization	(Optional) Disable synchronization between BGP and an IGP.
Step 8	no auto-summary	(Optional) Disable automatic network summarization. By default, when a subnet is redistributed from an IGP into BGP, only the network route is inserted into the BGP table.

	Command	Purpose
Step 9	bgp fast-external-fallover	(Optional) Automatically reset a BGP session when a link between external neighbors goes down. By default, the session is not immediately reset.
Step 10	end	Return to privileged EXEC mode.
Step 11	show ip bgp network <i>network-number</i> show ip bgp neighbor	Verify the configuration.
Step 12	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no router bgp** *autonomous-system* global configuration command to remove a BGP AS. Use the **no network** *network-number* router configuration command to remove the network from the BGP table. Use the **no neighbor** {*ip-address* | *peer-group-name*} **remote-as** *number* router configuration command to remove a neighbor. Use the **no neighbor** {*ip-address* | *peer-group-name*} **remove-private-as** router configuration command to include private AS numbers in updates to a neighbor. Use the **synchronization** router configuration command to re-enable synchronization.

These examples show how to configure BGP on the routers in [Figure 28-5](#).

Router A:

```
Switch(config)# router bgp 100
Switch(config-router)# neighbor 129.213.1.1 remote-as 200
```

Router B:

```
Switch(config)# router bgp 200
Switch(config-router)# neighbor 129.213.1.2 remote-as 100
Switch(config-router)# neighbor 175.220.1.2 remote-as 200
```

Router C:

```
Switch(config)# router bgp 200
Switch(config-router)# neighbor 175.220.212.1 remote-as 200
Switch(config-router)# neighbor 192.208.10.1 remote-as 300
```

Router D:

```
Switch(config)# router bgp 300
Switch(config-router)# neighbor 192.208.10.2 remote-as 200
```

To verify that BGP peers are running, use the **show ip bgp neighbors** privileged EXEC command. This is the output of this command on Router A:

```
Switch# show ip bgp neighbors
```

```
BGP neighbor is 129.213.1.1, remote AS 200, external link
BGP version 4, remote router ID 175.220.212.1
BGP state = established, table version = 3, up for 0:10:59
Last read 0:00:29, hold time is 180, keepalive interval is 60 seconds
Minimum time between advertisement runs is 30 seconds
Received 2828 messages, 0 notifications, 0 in queue
Sent 2826 messages, 0 notifications, 0 in queue
Connections established 11; dropped 10
```

Anything other than *state = established* means that the peers are not running. The remote router ID is the highest IP address on that router (or the highest loopback interface). Each time the table is updated with new information, the table version number increments. A table version number that continually increments means that a route is flapping, causing continual routing updates.

For exterior protocols, a reference to an IP network from the **network** router configuration command controls only which networks are advertised. This is in contrast to Interior Gateway Protocols (IGPs), such as IGRP, which also use the **network** command to determine where to send updates.

For detailed descriptions of BGP configuration, refer to the “Configuring BGP” chapter in the *Cisco IOS IP and IP Routing Configuration Guide*. For details about specific commands, refer to the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*. See [Appendix C, “Unsupported Commands in Cisco IOS Release 12.1\(14\)AX,”](#) for a list of BGP commands that are visible but not supported by the switch.

Managing Routing Policy Changes

Routing policies for a peer include all the configurations that might affect inbound or outbound routing table updates. When you have defined two routers as BGP neighbors, they form a BGP connection and exchange routing information. If you later change a BGP filter, weight, distance, version, or timer, or make a similar configuration change, you must reset the BGP sessions so that the configuration changes take effect.

There are two types of reset, hard reset and soft reset. Cisco IOS software releases 12.1 and later support a soft reset without any prior configuration. To use a soft reset without preconfiguration, both BGP peers must support the soft route refresh capability, which is advertised in the OPEN message sent when the peers establish a TCP session. A soft reset allows the dynamic exchange of route refresh requests and routing information between BGP routers and the subsequent re-advertisement of the respective outbound routing table.

- When soft reset generates inbound updates from a neighbor, it is called *dynamic inbound soft reset*.
- When soft reset sends a set of updates to a neighbor, it is called *outbound soft reset*.

A soft inbound reset causes the new inbound policy to take effect. A soft outbound reset causes the new local outbound policy to take effect without resetting the BGP session. As a new set of updates is sent during outbound policy reset, a new inbound policy can also take effect.

[Table 28-11](#) lists the advantages and disadvantages hard reset and soft reset.

Table 28-11 Advantages and Disadvantages of Hard and Soft Resets

Type of Reset	Advantages	Disadvantages
Hard reset	No memory overhead	The prefixes in the BGP, IP, and FIB tables provided by the neighbor are lost. Not recommended.
Outbound soft reset	No configuration, no storing of routing table updates	Does not reset inbound routing table updates.
Dynamic inbound soft reset	Does not clear the BGP session and cache Does not require storing of routing table updates and has no memory overhead	Both BGP routers must support the route refresh capability (in Cisco IOS Release 12.1 and later).

Beginning in privileged EXEC mode, follow these steps to determine if a BGP peer supports the route refresh capability and to reset the BGP session:

	Command	Purpose
Step 1	show ip bgp neighbors	Display whether a neighbor supports the route refresh capability. When supported, this message appears for the router: <i>Received route refresh capability from peer.</i>
Step 2	clear ip bgp { * <i>address</i> <i>peer-group-name</i> }	Reset the routing table on the specified connection. • Enter an asterisk (*) to specify that all connections be reset. • Enter an IP <i>address</i> to specify the connection to be reset. • Enter a peer group name to reset the peer group.
Step 3	clear ip bgp { * <i>address</i> <i>peer-group-name</i> } soft out	(Optional) Perform an outbound soft reset to reset the inbound routing table on the specified connection. Use this command if route refresh is supported. • Enter an asterisk (*) to specify that all connections be reset. • Enter an IP <i>address</i> to specify the connection to be reset. • Enter a peer group name to reset the peer group.
Step 4	show ip bgp show ip bgp neighbors	Verify the reset by checking information about the routing table and about BGP neighbors.

Configuring BGP Decision Attributes

When a BGP speaker receives updates from multiple autonomous systems that describe different paths to the same destination, it must choose the single best path for reaching that destination. When chosen, the selected path is entered into the BGP routing table and propagated to its neighbors. The decision is based on the value of attributes that the update contains and other BGP-configurable factors.

When a BGP peer learns two EBGp paths for a prefix from a neighboring AS, it chooses the best path and inserts that path in the IP routing table. If BGP multipath support is enabled and the EBGp paths are learned from the same neighboring autonomous systems, instead of a single best path, multiple paths are installed in the IP routing table. Then, during packet switching, per-packet or per-destination load balancing is performed among the multiple paths. The **maximum-paths** router configuration command controls the number of paths allowed.

These factors summarize the order in which BGP evaluates the attributes for choosing the best path:

1. If the path specifies a next hop that is inaccessible, drop the update. The BGP next-hop attribute, automatically determined by the software, is the IP address of the next hop that is going to be used to reach a destination. For EBGp, this is usually the IP address of the neighbor specified by the **neighbor remote-as** router configuration command. You can disable next-hop processing by using route maps or the **neighbor next-hop-self** router configuration command.
2. Prefer the path with the largest weight (a Cisco proprietary parameter). The weight attribute is local to the router and not propagated in routing updates. By default, the weight attribute is 32768 for paths that the router originates and zero for other paths. Routes with the largest weight are preferred. You can use access lists, route maps, or the **neighbor weight** router configuration command to set weights.

3. Prefer the route with the highest local preference. Local preference is part of the routing update and exchanged among routers in the same AS. The default value of the local preference attribute is 100. You can set local preference by using the **bgp default local-preference** router configuration command or by using a route map.
4. Prefer the route that was originated by BGP running on the local router.
5. Prefer the route with the shortest AS path.
6. Prefer the route with the lowest origin type. An interior route or IGP is lower than a route learned by EGP, and an EGP-learned route is lower than one of unknown origin or learned in another way.
7. Prefer the route with the lowest multi-exit discriminator (MED) metric attribute if the neighboring AS is the same for all routes considered. You can configure the MED by using route maps or by using the **default-metric** router configuration command. When an update is sent to an IBGP peer, the MED is included.
8. Prefer the external (EBGP) path over the internal (IBGP) path.
9. Prefer the route that can be reached through the closest IGP neighbor (the lowest IGP metric). This means that the router will prefer the shortest internal path within the AS to reach the destination (the shortest path to the BGP next-hop).
10. If the following conditions are all true, insert the route for this path into the IP routing table:
 - Both the best route and this route are external.
 - Both the best route and this route are from the same neighboring autonomous system.
 - maximum-paths is enabled.
11. If multipath is not enabled, prefer the route with the lowest IP address value for the BGP router ID. The router ID is usually the highest IP address on the router or the loopback (virtual) address, but might be implementation-specific.

Beginning in privileged EXEC mode, follow these steps to configure some decision attributes:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system</i>	Enable a BGP routing process, assign it an AS number, and enter router configuration mode.
Step 3	bgp best-path as-path ignore	(Optional) Configure the router to ignore AS path length in selecting a route.
Step 4	neighbor { <i>ip-address</i> <i>peer-group-name</i> } next-hop-self	(Optional) Disable next-hop processing on BGP updates to a neighbor by entering a specific IP address to be used instead of the next-hop address.
Step 5	neighbor { <i>ip-address</i> <i>peer-group-name</i> } weight <i>weight</i>	(Optional) Assign a weight to a neighbor connection. Acceptable values are from 0 to 65535; the largest weight is the preferred route. Routes learned through another BGP peer have a default weight of 0; routes sourced by the local router have a default weight of 32768.
Step 6	default-metric <i>number</i>	(Optional) Set a MED metric to set preferred paths to external neighbors. All routes without a MED will also be set to this value. The range is 1 to 4294967295. The lowest value is the most desirable.

	Command	Purpose
Step 7	bgp bestpath med missing-as-worst	(Optional) Configure the switch to consider a missing MED as having a value of infinity, making the path without a MED value the least desirable path.
Step 8	bgp always-compare med	(Optional) Configure the switch to compare MEDs for paths from neighbors in different autonomous systems. By default, MED comparison is only done among paths in the same AS.
Step 9	bgp bestpath med confed	(Optional) Configure the switch to consider the MED in choosing a path from among those advertised by different subautonomous systems within a confederation.
Step 10	bgp deterministic med	(Optional) Configure the switch to consider the MED variable when choosing among routes advertised by different peers in the same AS.
Step 11	bgp default local-preference <i>value</i>	(Optional) Change the default local preference value. The range is 0 to 4294967295; the default value is 100. The highest local preference value is preferred.
Step 12	maximum-paths <i>number</i>	(Optional) Configure the number of paths to be added to the IP routing table. The default is to only enter the best path in the routing table. The range is from 1 to 8. Having multiple paths allows load balancing among the paths.
Step 13	end	Return to privileged EXEC mode.
Step 14	show ip bgp show ip bgp neighbors	Verify the reset by checking information about the routing table and about BGP neighbors.
Step 15	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no** form of each command to return to the default state.

Configuring BGP Filtering with Route Maps

Within BGP, route maps can be used to control and to modify routing information and to define the conditions by which routes are redistributed between routing domains. See the [“Using Route Maps to Redistribute Routing Information” section on page 28-90](#) for more information about route maps. Each route map has a name that identifies the route map (*map tag*) and an optional sequence number.

Beginning in privileged EXEC mode, follow these steps to use a route map to disable next-hop processing:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	route-map <i>map-tag</i> [[permit deny] <i>sequence-number</i>]]	Create a route map, and enter route-map configuration mode.

	Command	Purpose
Step 3	set ip next-hop <i>ip-address</i> [... <i>ip-address</i>] [<i>peer-address</i>]	(Optional) Set a route map to disable next-hop processing - In an inbound route map, set the next hop of matching routes to be the neighbor peering address, overriding third-party next hops. - In an outbound route map of a BGP peer, set the next hop to the peering address of the local router, disabling the next-hop calculation.
Step 4	end	Return to privileged EXEC mode.
Step 5	show route-map [<i>map-name</i>]	Display all route maps configured or only the one specified to verify configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no route-map** *map-tag* command to delete the route map. Use the **no set ip next-hop** *ip-address* command to re-enable next-hop processing.

Configuring BGP Filtering by Neighbor

You can filter BGP advertisements by using AS-path filters, such as the **as-path access-list** global configuration command and the **neighbor filter-list** router configuration command. You can also use access lists with the **neighbor distribute-list** router configuration command. Distribute-list filters are applied to network numbers. See the [“Controlling Advertising and Processing in Routing Updates” section on page 28-97](#) for information about the **distribute-list** command.

You can use route maps on a per-neighbor basis to filter updates and to modify various attributes. A route map can be applied to either inbound or outbound updates. Only the routes that pass the route map are sent or accepted in updates. On both inbound and outbound updates, matching is supported based on AS path, community, and network numbers. Autonomous system path matching requires the **match as-path access-list** route-map command, community based matching requires the **match community-list** route-map command, and network-based matching requires the **ip access-list** global configuration command.

Beginning in privileged EXEC mode, follow these steps to apply a per-neighbor route map:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system</i>	Enable a BGP routing process, assign it an AS number, and enter router configuration mode.
Step 3	neighbor { <i>ip-address</i> <i>peer-group name</i> } distribute-list { <i>access-list-number</i> <i>name</i> } { in out }	(Optional) Filter BGP routing updates to or from neighbors as specified in an access list. Note You can also use the neighbor prefix-list router configuration command to filter updates, but you cannot use both commands to configure the same BGP peer.
Step 4	neighbor { <i>ip-address</i> <i>peer-group name</i> } route-map <i>map-tag</i> { in out }	(Optional) Apply a route map to filter an incoming or outgoing route.
Step 5	end	Return to privileged EXEC mode.

	Command	Purpose
Step 6	show ip bgp neighbors	Verify the configuration.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no neighbor distribute-list** command to remove the access list from the neighbor. Use the **no neighbor route-map map-tag** router configuration command to remove the route map from the neighbor.

Another method of filtering is to specify an access list filter on both incoming and outbound updates, based on the BGP autonomous system paths. Each filter is an access list based on regular expressions. (Refer to the “Regular Expressions” appendix in the *Cisco IOS Dial Services Command Reference* for more information on forming regular expressions.) To use this method, define an autonomous system path access list, and apply it to updates to and from particular neighbors.

Beginning in privileged EXEC mode, follow these steps to configure BGP path filtering:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip as-path access-list access-list-number {permit deny} as-regular-expressions	Define a BGP-related access list.
Step 3	router bgp autonomous-system	Enter BGP router configuration mode.
Step 4	neighbor {ip-address peer-group name} filter-list {access-list-number name} {in out weight weight}	Establish a BGP filter based on an access list.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip bgp neighbors [paths regular-expression]	Verify the configuration.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Configuring Prefix Lists for BGP Filtering

You can use prefix lists as an alternative to access lists in many BGP route filtering commands, including the **neighbor distribute-list** router configuration command. The advantages of using prefix lists include performance improvements in loading and lookup of large lists, incremental update support, easier CLI configuration, and greater flexibility.

Filtering by a prefix list involves matching the prefixes of routes with those listed in the prefix list, as when matching access lists. When there is a match, the route is used. Whether a prefix is permitted or denied is based upon these rules:

- An empty prefix list permits all prefixes.
- An implicit deny is assumed if a given prefix does not match any entries in a prefix list.
- When multiple entries of a prefix list match a given prefix, the sequence number of a prefix list entry identifies the entry with the lowest sequence number.

By default, sequence numbers are generated automatically and incremented in units of five. If you disable the automatic generation of sequence numbers, you must specify the sequence number for each entry. You can specify sequence values in any increment. If you specify increments of one, you cannot insert additional entries into the list; if you choose very large increments, you might run out of values.

You do not need to specify a sequence number when removing a configuration entry. **Show** commands include the sequence numbers in their output.

Before using a prefix list in a command, you must set up the prefix list. Beginning in privileged EXEC mode, follow these steps to create a prefix list or to add an entry to a prefix list:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip prefix-list <i>list-name</i> [seq <i>seq-value</i>] deny permit <i>network/len</i> [ge <i>ge-value</i>] [le <i>le-value</i>]	Create a prefix list with an optional sequence number to deny or permit access for matching conditions. You must enter at least one permit or deny clause. <i>network/len</i> is the network number and length (in bits) of the network mask. (Optional) ge and le values specify the range of the prefix length to be matched. The specified <i>ge-value</i> and <i>le-value</i> must satisfy this condition: $len < ge\text{-}value < le\text{-}value < 32$
Step 3	ip prefix-list <i>list-name</i> seq <i>seq-value</i> deny permit <i>network/len</i> [ge <i>ge-value</i>] [le <i>le-value</i>]	(Optional) Add an entry to a prefix list, and assign a sequence number to the entry.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip prefix list [detail summary] <i>name</i> [<i>network/len</i>] [seq <i>seq-num</i>] [longer] [first-match]	Verify the configuration by displaying information about a prefix list or prefix list entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete a prefix list and all of its entries, use the **no ip prefix-list** *list-name* global configuration command. To delete an entry from a prefix list, use the **no ip prefix-list seq** *seq-value* global configuration command. To disable automatic generation of sequence numbers, use the **no ip prefix-list sequence number** command; to reenable automatic generation, use the **ip prefix-list sequence number** command. To clear the hit-count table of prefix list entries, use the **clear ip prefix-list** privileged EXEC command.

Configuring BGP Community Filtering

One way that BGP controls the distribution of routing information based on the value of the COMMUNITIES attribute. The attribute is a way to groups destinations into communities and to apply routing decisions based on the communities. This method simplifies configuration of a BGP speaker to control distribution of routing information.

A *community* is a group of destinations that share some common attribute. Each destination can belong to multiple communities. AS administrators can define to which communities a destination belongs. By default, all destinations belong to the general Internet community. The community is identified by the COMMUNITIES attribute, an optional, transitive, global attribute in the numerical range from 1 to 4294967200. These are some predefined, well-known communities:

- **internet**—Advertise this route to the Internet community. All routers belong to it.
- **no-export**—Do not advertise this route to EBGp peers.
- **no-advertise**—Do not advertise this route to any peer (internal or external).
- **local-as**—Do not advertise this route to peers outside the local autonomous system.

Based on the community, you can control which routing information to accept, prefer, or distribute to other neighbors. A BGP speaker can set, append, or modify the community of a route when learning, advertising, or redistributing routes. When routes are aggregated, the resulting aggregate has a COMMUNITIES attribute that contains all communities from all the initial routes.

You can use community lists to create groups of communities to use in a match clause of a route map. As with an access list, a series of community lists can be created. Statements are checked until a match is found. As soon as one statement is satisfied, the test is concluded.

To set the COMMUNITIES attribute and match clauses based on communities, see the **match community-list** and **set community** route-map configuration commands in the [“Using Route Maps to Redistribute Routing Information”](#) section on page 28-90.

By default, no COMMUNITIES attribute is sent to a neighbor. You can specify that the COMMUNITIES attribute be sent to the neighbor at an IP address by using the **neighbor send-community** router configuration command.

Beginning in privileged EXEC mode, follow these steps to create and to apply a community list:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip community-list <i>community-list-number</i> { permit deny } <i>community-number</i>	Create a community list, and assign it a number. - The <i>community-list-number</i> is an integer from 1 to 99 that identifies one or more permit or deny groups of communities. - The <i>community-number</i> is the number configured by a set community route-map configuration command.
Step 3	router bgp <i>autonomous-system</i>	Enter BGP router configuration mode.
Step 4	neighbor { <i>ip-address</i> <i>peer-group name</i> } send-community	Specify that the COMMUNITIES attribute be sent to the neighbor at this IP address.
Step 5	set comm-list <i>list-num</i> delete	(Optional) Remove communities from the community attribute of an inbound or outbound update that match a standard or extended community list specified by a route map.
Step 6	exit	Return to global configuration mode.
Step 7	ip bgp-community new-format	(Optional) Display and parse BGP communities in the format AA:NN. A BGP community is displayed in a two-part format 2 bytes long. The Cisco default community format is in the format NNAA. In the most recent RFC for BGP, a community takes the form AA:NN, where the first part is the AS number and the second part is a 2-byte number.
Step 8	end	Return to privileged EXEC mode.
Step 9	show ip bgp community	Verify the configuration.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Configuring BGP Neighbors and Peer Groups

Often many BGP neighbors are configured with the same update policies (that is, the same outbound route maps, distribute lists, filter lists, update source, and so on). Neighbors with the same update policies can be grouped into peer groups to simplify configuration and to make updating more efficient. When you have configured many peers, we recommend this approach.

To configure a BGP peer group, you create the peer group, assign options to the peer group, and add neighbors as peer group members. You configure the peer group by using the **neighbor** router configuration commands. By default, peer group members inherit all the configuration options of the peer group, including the remote-as (if configured), version, update-source, out-route-map, out-filter-list, out-dist-list, minimum-advertisement-interval, and next-hop-self. All peer group members also inherit changes made to the peer group. Members can also be configured to override the options that do not affect outbound updates.

To assign configuration options to an individual neighbor, specify any of these router configuration commands by using the neighbor IP address. To assign the options to a peer group, specify any of the commands by using the peer group name. You can disable a BGP peer or peer group without removing all the configuration information by using the **neighbor shutdown** router configuration command.

Beginning in privileged EXEC mode, use these commands to configure BGP peers:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system</i>	Enter BGP router configuration mode.
Step 3	neighbor <i>peer-group-name</i> peer-group	Create a BGP peer group.
Step 4	neighbor <i>ip-address</i> peer-group <i>peer-group-name</i>	Make a BGP neighbor a member of the peer group.
Step 5	neighbor { <i>ip-address</i> <i>peer-group-name</i> } remote-as <i>number</i>	Specify a BGP neighbor. If a peer group is not configured with a remote-as <i>number</i> , use this command to create peer groups containing EBGP neighbors. The range is 1 to 65535.
Step 6	neighbor { <i>ip-address</i> <i>peer-group-name</i> } description <i>text</i>	(Optional) Associate a description with a neighbor.
Step 7	neighbor { <i>ip-address</i> <i>peer-group-name</i> } default-originate [route-map <i>map-name</i>]	(Optional) Allow a BGP speaker (the local router) to send the default route 0.0.0.0 to a neighbor for use as a default route.
Step 8	neighbor { <i>ip-address</i> <i>peer-group-name</i> } send-community	(Optional) Specify that the COMMUNITIES attribute be sent to the neighbor at this IP address.
Step 9	neighbor { <i>ip-address</i> <i>peer-group-name</i> } update-source <i>interface</i>	(Optional) Allow internal BGP sessions to use any operational interface for TCP connections.
Step 10	neighbor { <i>ip-address</i> <i>peer-group-name</i> } ebgp-multihop	(Optional) Allow BGP sessions, even when the neighbor is not on a directly connected segment. The multihop session is not established if the only route to the multihop peer's address is the default route (0.0.0.0).
Step 11	neighbor { <i>ip-address</i> <i>peer-group-name</i> } local-as <i>number</i>	(Optional) Specify an AS number to use as the local AS. The range is 1 to 65535.
Step 12	neighbor { <i>ip-address</i> <i>peer-group-name</i> } advertisement-interval <i>seconds</i>	(Optional) Set the minimum interval between sending BGP routing updates.
Step 13	neighbor { <i>ip-address</i> <i>peer-group-name</i> } maximum-prefix <i>maximum</i> [<i>threshold</i>]	(Optional) Control how many prefixes can be received from a neighbor. The range is 1 to 4294967295. The <i>threshold</i> (optional) is the percentage of maximum at which a warning message is generated. The default is 75 percent.
Step 14	neighbor { <i>ip-address</i> <i>peer-group-name</i> } next-hop-self	(Optional) Disable next-hop processing on the BGP updates to a neighbor.

	Command	Purpose
Step 15	neighbor { <i>ip-address</i> <i>peer-group-name</i> } password <i>string</i>	(Optional) Set MD5 authentication on a TCP connection to a BGP peer. The same password must be configured on both BGP peers, or the connection between them is not made.
Step 16	neighbor { <i>ip-address</i> <i>peer-group-name</i> } route-map <i>map-name</i> { in out }	(Optional) Apply a route map to incoming or outgoing routes.
Step 17	neighbor { <i>ip-address</i> <i>peer-group-name</i> } send-community	(Optional) Specify that the COMMUNITIES attribute be sent to the neighbor at this IP address.
Step 18	neighbor { <i>ip-address</i> <i>peer-group-name</i> } timers <i>keepalive holdtime</i>	(Optional) Set timers for the neighbor or peer group. - The <i>keepalive</i> interval is the time within which keepalive messages are sent to peers. The range is 1 to 4294967295 seconds; the default is 60. - The <i>holdtime</i> is the interval after which a peer is declared inactive after not receiving a keepalive message from it. The range is 1 to 4294967295 seconds; the default is 180.
Step 19	neighbor { <i>ip-address</i> <i>peer-group-name</i> } weight <i>weight</i>	(Optional) Specify a weight for all routes from a neighbor.
Step 20	neighbor { <i>ip-address</i> <i>peer-group-name</i> } distribute-list { <i>access-list-number</i> <i>name</i> } { in out }	(Optional) Filter BGP routing updates to or from neighbors, as specified in an access list.
Step 21	neighbor { <i>ip-address</i> <i>peer-group-name</i> } filter-list <i>access-list-number</i> { in out weight <i>weight</i> }	(Optional) Establish a BGP filter.
Step 22	neighbor { <i>ip-address</i> <i>peer-group-name</i> } version <i>value</i>	(Optional) Specify the BGP version to use when communicating with a neighbor.
Step 23	neighbor { <i>ip-address</i> <i>peer-group-name</i> } soft-reconfiguration inbound	(Optional) Configure the software to start storing received updates.
Step 24	end	Return to privileged EXEC mode.
Step 25	show ip bgp neighbors	Verify the configuration.
Step 26	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable an existing BGP neighbor or neighbor peer group, use the **neighbor shutdown** router configuration command. To enable a previously existing neighbor or neighbor peer group that had been disabled, use the **no neighbor shutdown** router configuration command.

Configuring Aggregate Addresses

Classless interdomain routing (CIDR) enables you to create aggregate routes (or *supernets*) to minimize the size of routing tables. You can configure aggregate routes in BGP either by redistributing an aggregate route into BGP or by creating an aggregate entry in the BGP routing table. An aggregate address is added to the BGP table when there is at least one more specific entry in the BGP table.

Beginning in privileged EXEC mode, use these commands to create an aggregate address in the routing table:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system</i>	Enter BGP router configuration mode.
Step 3	aggregate-address <i>address mask</i>	Create an aggregate entry in the BGP routing table. The aggregate route is advertised as coming from the AS, and the atomic aggregate attribute is set to indicate that information might be missing.
Step 4	aggregate-address <i>address mask as-set</i>	(Optional) Generate AS set path information. This command creates an aggregate entry following the same rules as the previous command, but the advertised path will be an AS_SET consisting of all elements contained in all paths. Do not use this keyword when aggregating many paths because this route must be continually withdrawn and updated.
Step 5	aggregate-address <i>address-mask summary-only</i>	(Optional) Advertise summary addresses only.
Step 6	aggregate-address <i>address mask suppress-map</i> <i>map-name</i>	(Optional) Suppress selected, more specific routes.
Step 7	aggregate-address <i>address mask advertise-map</i> <i>map-name</i>	(Optional) Generate an aggregate based on conditions specified by the route map.
Step 8	aggregate-address <i>address mask attribute-map</i> <i>map-name</i>	(Optional) Generate an aggregate with attributes specified in the route map.
Step 9	end	Return to privileged EXEC mode.
Step 10	show ip bgp neighbors [advertised-routes]	Verify the configuration.
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete an aggregate entry, use the **no aggregate-address** *address mask* router configuration command. To return options to the default values, use the command with keywords.

Configuring Routing Domain Confederations

One way to reduce the IBGP mesh is to divide an autonomous system into multiple subautonomous systems and to group them into a single confederation that appears as a single autonomous system. Each autonomous system is fully meshed within itself and has a few connections to other autonomous systems in the same confederation. Even though the peers in different autonomous systems have EBGP sessions, they exchange routing information as if they were IBGP peers. Specifically, the next hop, MED, and local preference information is preserved. You can then use a single IGP for all of the autonomous systems.

To configure a BGP confederation, you must specify a confederation identifier that acts as the autonomous system number for the group of autonomous systems.

Beginning in privileged EXEC mode, use these commands to configure a BGP confederation:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system</i>	Enter BGP router configuration mode.

	Command	Purpose
Step 3	bgp confederation identifier <i>autonomous-system</i>	Configure a BGP confederation identifier.
Step 4	bgp confederation peers <i>autonomous-system</i> [<i>autonomous-system ...</i>]	Specify the autonomous systems that belong to the confederation and that will be treated as special EBGp peers.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip bgp neighbor show ip bgp network	Verify the configuration.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Configuring BGP Route Reflectors

BGP requires that all of the IBGP speakers be fully meshed. When a router receives a route from an external neighbor, it must advertise it to all internal neighbors. To prevent a routing information loop, all IBGP speakers must be connected. The internal neighbors do not send routes learned from internal neighbors to other internal neighbors.

With route reflectors, all IBGP speakers need not be fully meshed because another method is used to pass learned routes to neighbors. When you configure an internal BGP peer to be a *route reflector*, it is responsible for passing IBGP learned routes to a set of IBGP neighbors. The internal peers of the route reflector are divided into two groups: *client peers* and *nonclient peers* (all the other routers in the autonomous system). A route reflector reflects routes between these two groups. The route reflector and its client peers form a *cluster*. The nonclient peers must be fully meshed with each other, but the client peers need not be fully meshed. The clients in the cluster do not communicate with IBGP speakers outside their cluster.

When the route reflector receives an advertised route, it takes one of these actions, depending on the neighbor:

- A route from an external BGP speaker is advertised to all clients and nonclient peers.
- A route from a nonclient peer is advertised to all clients.
- A route from a client is advertised to all clients and nonclient peers. Hence, the clients need not be fully meshed.

Usually a cluster of clients have a single route reflector, and the cluster is identified by the route reflector router ID. To increase redundancy and to avoid a single point of failure, a cluster might have more than one route reflector. In this case, all route reflectors in the cluster must be configured with the same 4-byte cluster ID so that a route reflector can recognize updates from route reflectors in the same cluster. All the route reflectors serving a cluster should be fully meshed and should have identical sets of client and nonclient peers.

Beginning in privileged EXEC mode, use these commands to configure a route reflector and clients:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system</i>	Enter BGP router configuration mode.
Step 3	neighbor <i>ip-address</i> <i>peer-group-name</i> route-reflector-client	Configure the local router as a BGP route reflector and the specified neighbor as a client.
Step 4	bgp cluster-id <i>cluster-id</i>	(Optional) Configure the cluster ID if the cluster has more than one route reflector.

	Command	Purpose
Step 5	no bgp client-to-client reflection	(Optional) Disable client-to-client route reflection. By default, the routes from a route reflector client are reflected to other clients. However, if the clients are fully meshed, the route reflector does not need to reflect routes to clients.
Step 6	end	Return to privileged EXEC mode.
Step 7	show ip bgp	Verify the configuration. Display the originator ID and the cluster-list attributes.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Configuring Route Dampening

Route flap dampening is a BGP feature designed to minimize the propagation of flapping routes across an internetwork. A route is considered to be flapping when it is repeatedly available, then unavailable, then available, then unavailable, and so on. When route dampening is enabled, a numeric *penalty* value is assigned to a route when it flaps. When a route's accumulated penalties reach a configurable limit, BGP suppresses advertisements of the route, even if the route is running. The *reuse limit* is a configurable value that is compared with the penalty. If the penalty is less than the reuse limit, a suppressed route that is up is advertised again.

Dampening is not applied to routes that are learned by IBGP. This policy prevents the IBGP peers from having a higher penalty for routes external to the AS.

Beginning in privileged EXEC mode, use these commands to configure BGP route dampening:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system</i>	Enter BGP router configuration mode.
Step 3	bgp dampening	Enable BGP route dampening.
Step 4	bgp dampening <i>half-life reuse suppress max-suppress</i> [route-map <i>map</i>]	(Optional) Change the default values of route dampening factors.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip bgp flap-statistics [{ regex <i>regex</i> } { filter-list <i>list</i> } { <i>address mask</i> [longer-prefix]}]	(Optional) Monitor the flaps of all paths that are flapping. The statistics are deleted when the route is not suppressed and is stable.
Step 7	show ip bgp dampened-paths	(Optional) Display the dampened routes, including the time remaining before they are suppressed.
Step 8	clear ip bgp flap-statistics [{ regex <i>regex</i> } { filter-list <i>list</i> } { <i>address mask</i> [longer-prefix]}]	(Optional) Clear BGP flap statistics to make it less likely that a route will be dampened.
Step 9	clear ip bgp dampening	(Optional) Clear route dampening information, and unsuppress the suppressed routes.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable flap dampening, use the **no bgp dampening** router configuration command without keywords. To set dampening factors back to the default values, use the **no bgp dampening** router configuration command with values.

Monitoring and Maintaining BGP

You can remove all contents of a particular cache, table, or database. This might be necessary when the contents of the particular structure have become or are suspected to be invalid.

You can display specific statistics, such as the contents of BGP routing tables, caches, and databases. You can use the information to determine resource utilization and solve network problems. You can also display information about node reachability and discover the routing path your device's packets are taking through the network.

Table 28-9 lists the privileged EXEC commands for clearing and displaying BGP. For explanations of the display fields, refer to the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

Table 28-12 IP BGP Clear and Show Commands

Command	Purpose
clear ip bgp <i>address</i>	Reset a particular BGP connection.
clear ip bgp *	Reset all BGP connections.
clear ip bgp peer-group <i>tag</i>	Remove all members of a BGP peer group.
show ip bgp <i>prefix</i>	Display peer groups and peers not in peer groups to which the prefix has been advertised. Also display prefix attributes such as the next hop and the local prefix.
show ip bgp cidr-only	Display all BGP routes that contain subnet and supernet network masks.
show ip bgp community [<i>community-number</i>] [exact]	Display routes that belong to the specified communities.
show ip bgp community-list <i>community-list-number</i> [exact-match]	Display routes that are permitted by the community list.
show ip bgp filter-list <i>access-list-number</i>	Display routes that are matched by the specified AS path access list.
show ip bgp inconsistent-as	Display the routes with inconsistent originating autonomous systems.
show ip bgp regexp <i>regular-expression</i>	Display the routes that have an AS path that matches the specified regular expression entered on the command line.
show ip bgp	Display the contents of the BGP routing table.
show ip bgp neighbors [<i>address</i>]	Display detailed information on the BGP and TCP connections to individual neighbors.
show ip bgp neighbors [<i>address</i>] [advertised-routes dampened-routes flap-statistics paths <i>regular-expression</i> received-routes routes]	Display routes learned from a particular BGP neighbor.
show ip bgp paths	Display all BGP paths in the database.
show ip bgp peer-group [<i>tag</i>] [summary]	Display information about BGP peer groups.
show ip bgp summary	Display the status of all BGP connections.

You can also enable the logging of messages generated when a BGP neighbor resets, comes up, or goes down by using the **bgp log-neighbor changes** router configuration command.

Configuring ISO CLNS Routing

The International Organization for Standardization (ISO) Connectionless Network Service (CLNS) protocol is a standard for the network layer of the Open System Interconnection (OSI) model. Addresses in the ISO network architecture are referred to as network service access point (NSAP) addresses and network entity titles (NETs). Each node in an OSI network has one or more NETs. In addition, each node has many NSAP addresses.

When you enable connectionless routing on the switch by using the **clns routing** global configuration command, the switch makes only forwarding decisions, with no routing-related functionality. For dynamic routing, you must also enable a routing protocol. The switch supports these dynamic routing protocols for ISO CLNS networks:

- ISO Interior Gateway Routing Protocol (IGRP)
- Intermediate System-to-Intermediate System (IS-IS) routing

When dynamically routing, you can choose either ISO IGRP or IS-IS, or you can enable both routing protocols at the same time. Both routing protocols support the concept of *areas*. Within an area, all routers know how to reach all the system IDs. Between areas, routers know how to reach the proper area.

ISO IGRP supports three levels of routing: *system routing*, *area routing*, and *interdomain routing*. Routing across domains (interdomain routing) can be done either statically or dynamically with ISO IGRP. IS-IS supports two levels of routing: *station routing* (within an area) and *area routing* (between areas).

The key difference between the ISO IGRP and IS-IS NSAP addressing schemes is in the definition of area addresses. Both use the system ID for Level 1 routing (routing within an area). However, they differ in the way addresses are specified for area routing. An ISO IGRP NSAP address includes three separate fields for routing: the *domain*, *area*, and *system ID*. An IS-IS address includes two fields: a single continuous *area* field (comprising the domain and area fields) and the *system ID*.

This section contains this ISO CLNS configuration information:

- [Configuring ISO IGRP Dynamic Routing, page 28-63](#)
- [Configuring IS-IS Dynamic Routing, page 28-66](#)



Note

For more detailed information about ISO CLNS, refer to the *Cisco IOS Apollo Domain, Banyan VINES, DECnet, ISO CLNS and XNS Configuration Guide for Release 12.1*. For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco IOS Apollo Domain, Banyan VINES, DECnet, ISO CLNS and XNS Command Reference for Release 12.1*, use the IOS command reference master index, or search online.

Configuring ISO IGRP Dynamic Routing

The ISO IGRP is a dynamic distance-vector protocol designed by Cisco Systems for routing an autonomous system that contains large, arbitrarily complex networks with diverse bandwidth and delay characteristics. This section briefly describes how to configure ISO IGRP. It includes this information:

- [Default ISO IGRP Configuration, page 28-64](#)
- [Enabling ISO IGRP, page 28-64](#)
- [Configuring Optional ISO IGRP Parameters, page 28-65](#)

Default ISO IGRP Configuration

Table 28-5 shows the default ISO IGRP configuration.

Table 28-13 Default ISO IGRP Configuration

Feature	Default Setting
Metric weights	Type of service (ToS): 0 Metric constants: k1: 1, k2: 0, k3: 1, k4: 0, k5: 0
Timers	Update interval: 90 seconds. Invalid interval: 135 seconds. Holddown interval: 145 seconds.
Split horizon	Enabled.

Enabling ISO IGRP

To configure ISO IGRP dynamic routing, you must enable the ISO IGRP routing process, identify the address for the router, and specify the interfaces that are to route ISO IGRP. Optionally, you can set a level for your routing updates when you configure the interfaces. CLNS routing is enabled by default on routers when you configure ISO IGRP. You can specify up to ten ISO IGRP routing processes.

Beginning in privileged EXEC mode, follow these steps to enable ISO IGRP:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	clns routing	Enable ISO connectionless routing on the switch.
Step 3	router iso- igrp [<i>tag</i>]	Enable the ISO IGRP routing process and enter router configuration mode. (Optional) Assign a meaningful name (<i>tag</i>) for the routing process.
Step 4	net <i>network-entity-title</i>	Configure the NET or address for the routing process. You can configure only one NET per routing process. Enter the NET or process in the format XX.XXXX. XXX.XX
Step 5	exit	Return to global configuration mode.
Step 6	interface <i>interface-id</i>	Specify an interface to route ISO IGRP and enter interface configuration mode. If the interface is not already configured as a Layer 3 interface, enter the no switchport command to put it into Layer 3 mode.
Step 7	clns router iso-igrp [[<i>tag</i>] [level 2]]	Enable ISO IGRP on the specified interface. (Optional) <i>tag</i>—routing process name. (Optional) Set the routing type (level 2) for routing updates. When selected, this option reduces the amount of router-to-router traffic because the interface sends out only Level 2 routing updates and does not pass on Level 1 information.
Step 8	exit	Return to global configuration mode. Repeat steps 5 and 6 to enable ISO IGRP on other interfaces.
Step 9	end	Return to privileged EXEC mode.

	Command	Purpose
Step 10	show clns	Verify your entries.
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable ISO IGRP routing for the switch, use the **no router iso-igrp [tag]** global configuration command.

This example shows how to enable an ISO IGRP routing process defined as *test*, configure the net for the processing area 1, domain 47.0004.004d, enable the routing process on Gigabit Ethernet ports 1 and 2, and verify the configuration.

```
Switch(config)# clns routing
Switch(config)# router iso-igrp test
Switch(config-router)# net 47.0004.004d.0001.212223242526.00
Switch(config-router)# exit
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# clns router iso-igrp test
Switch(config-if)# exit
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# clns router iso-igrp test
Switch(config-if)# end
Switch# show clns
Global CLNS Information:
  2 Interfaces Enabled for CLNS
  NET: 37.0004.004d.0001.2122.2324.2526.00
  Configuration Timer: 60, Default Holding Timer: 300, Packet Lifetime 64
  ERPDUs requested on locally generated packets
  Intermediate system operation enabled (CLNS forwarding allowed)
  ISO-IGRP level-1 Router: test
    Routing for Domain: 37.0004.004D, Area: 0001
  ISO-IGRP level-2 Router: DOMAIN_test
    Routing for Domain: 37.0004.004D
```

Configuring Optional ISO IGRP Parameters

You can customize certain ISO IGRP parameters, such as timers and split horizon. You can adjust the metric constants used in the ISO IGRP metric calculation of reliability and load for specific applications.



Note

Adjusting the ISO IGRP metrics can dramatically affect network performance, so use care when making metric adjustments. Because of the complexity of this task, Cisco does not recommend changing the metrics without guidance from an experienced system designer.

You can also configure the bandwidth and delay associated with an interface to change the metrics. Refer to the *Cisco IOS Interface Command Reference* publication for details about the **bandwidth** and **delay** interface configuration commands used to set these metrics. Using these commands to change the values of the ISO IGRP metrics also changes the values of IP IGRP metrics.

You can adjust the basic timing parameters for ISO IGRP. The timers must be the same for all routers in the network.

Split horizon blocks information about routes from being advertised out the interface from which the information originated. This optimizes communication among multiple routers when links are broken. You can enable or disable split horizon on an interface.

Beginning in privileged EXEC mode, follow these steps to configure IGRP. Configuring the routing process is required; other steps are optional:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	clns routing	Enable ISO connectionless routing on the switch.
Step 3	router iso- igrp [<i>tag</i>]	Enable the ISO IGRP routing, and enter router configuration mode.
Step 4	metric weights <i>tos k1 k2 k3 k4 k5</i>	(Optional) Adjust the ISO IGRP metric. For TOS, only 0 is supported. The <i>k</i> values are metric constants used in the ISO IGRP equation that converts an IGRP metric vector into a scalar quantity. For <i>k1</i> - <i>k5</i> , the range is 0 to 4294967295. The defaults for <i>tos</i> , <i>k2</i> , <i>k4</i> , and <i>k5</i> are 0; for <i>k1</i> and <i>k3</i> , the default is 1. Higher numbers mean a greater multiplier effect.
Step 5	timers basic <i>update-interval</i> <i>invalid-interval holddown-interval</i>	(Optional) Adjust the ISO IGRP timers. The range for all timers is from 0 to 4294967295 seconds. <i>update-interval</i>—The time between sending routing updates. The default is 90 seconds. <i>invalid-interval</i>—The timer interval after which a route is declared invalid. The default is 135 seconds. <i>holddown-interval</i>—The time during which routing information about better paths is suppressed. The default is 145 seconds.
Step 6	interface <i>interface-id</i>	Specify an interface to route ISO IGRP and enter interface configuration mode. If the interface is not already configured as a Layer 3 interface, enter the no switchport command to put it into Layer 3 mode.
Step 7	[no] clns split-horizon	(Optional) Enable or disable split horizon for ISO IGRP updates on the interface. The default is enabled.
Step 8	end	Return to privileged EXEC mode.
Step 9	show clns [<i>interface</i>] protocol	Verify your entries.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the five (*k*) metric weights to their default values, use the **no metric weights** router configuration command. To restore the default timer values, use the **no timers basic** *update-interval invalid-interval holddown-interval* router configuration command.

Configuring IS-IS Dynamic Routing

IS-IS is an ISO dynamic routing protocol (described in ISO 105890). Unlike other routing protocols, enabling IS-IS requires that you create an IS-IS routing process and assign it to a specific interface, rather than to a network. You can specify more than one IS-IS routing process per Layer 3 switch or router by using the multiarea IS-IS configuration syntax. You then configure the parameters for each instance of the IS-IS routing process.

Small IS-IS networks are built as a single area that includes all the routers in the network. As the network grows larger, it is usually reorganized into a backbone area made up of the connected set of all Level 2 routers from all areas, which is in turn connected to local areas. Within a local area, routers know how to reach all system IDs. Between areas, routers know how to reach the backbone, and the backbone routers know how to reach other areas.

Routers establish Level 1 adjacencies to perform routing within a local area (station routing). Routers establish Level 2 adjacencies to perform routing between Level 1 areas (area routing).

A single Cisco router can participate in routing in up to 29 areas and can perform Level 2 routing in the backbone. In general, each routing process corresponds to an area. By default, the first instance of the routing process configured performs both Level 1 and Level 2 routing. You can configure additional router instances, which are automatically treated as Level 1 areas. You must configure the parameters for each instance of the IS-IS routing process individually.

For IS-IS multiarea routing, you can configure only one process to perform Level 2 routing, although you can define up to 29 Level 1 areas for each Cisco unit. If Level 2 routing is configured on any process, all additional processes are automatically configured as Level 1. You can configure this process to perform Level 1 routing at the same time. If Level 2 routing is not desired for a router instance, remove the Level 2 capability using the **is-type** global configuration command. Use the **is-type** command also to configure a different router instance as a Level 2 router.



Note

For more detailed information about IS-IS, refer to the “IP Routing Protocols” chapter of the *Cisco IOS IP and IP Routing Configuration Guide for Release 12.1*. For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

This section briefly describes how to configure IS-IS routing. It includes this information:

- [Default IS-IS Configuration, page 28-67](#)
- [Enabling IS-IS Routing, page 28-68](#)
- [Configuring IS-IS Global Parameters, page 28-70](#)
- [Configuring IS-IS Interface Parameters, page 28-72](#)

Default IS-IS Configuration

Table 28-14 shows the default IS-IS configuration.

Table 28-14 Default IS-IS Configuration

Feature	Default Setting
Ignore link-state PDU (LSP) errors	Enabled.
IS-IS type	Conventional IS-IS: the router acts as both a Level 1 (station) and a Level 2 (area) router. Multiarea IS-IS: the first instance of the IS-IS routing process is a Level 1-2 router. Remaining instances are Level 1 routers.
Default-information originate	Disabled.
Log IS-IS adjacency state changes.	Disabled.
LSP generation throttling timers	Maximum interval between two consecutive occurrences: 5 seconds. Initial LSP generation delay: 50 ms. Hold time between the first and second LSP generation: 5000 ms.
LSP maximum lifetime (without a refresh)	1200 seconds (20 minutes) before the LSP packet is deleted.
LSP refresh interval	Send LSP refreshes every 900 seconds (15 minutes).

Table 28-14 Default IS-IS Configuration (continued)

Feature	Default Setting
Maximum LSP packet size	1497 bytes.
Partial route computation (PRC) throttling timers	Maximum PRC wait interval: 5 seconds. Initial PRC calculation delay after a topology change: 2000 ms. Hold time between the first and second PRC calculation: 5000 ms.
Partition avoidance	Disabled.
Password	No area or domain password is defined, and authentication is disabled.
Set-overload-bit	Disabled. When enabled, if no arguments are entered, the overload bit is set immediately and remains set until you enter the no set-overload-bit command.
Shortest path first (SPF) throttling timers	Maximum interval between consecutive SFPS: 10 seconds. Initial SPF calculation after a topology change: 5500 ms. Holdtime between the first and second SPF calculation: 5500 ms.
Summary-address	Disabled.

Enabling IS-IS Routing

To enable IS-IS, you specify a name and NET for each routing process. You then enable IS-IS routing on the interface and specify the area for each instance of the routing process.

Beginning in privileged EXEC mode, follow these steps to enable IS-IS and specify the area for each instance of the IS-IS routing process:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	clsns routing	Enable ISO connectionless routing on the switch.
Step 3	router isis [<i>area tag</i>]	Enable the IS-IS routing for the specified routing process and enter IS-IS routing configuration mode. (Optional) Use the <i>area tag</i> argument to identify the area to which the IS-IS router is assigned. You must enter a value if you are configuring multiple IS-IS areas. The first IS-IS instance configured is Level 1-2 by default. Later instances are automatically Level 1. You can change the level of routing by using the is-type global configuration command.
Step 4	net <i>network-entity-title</i>	Configure the NETs for the routing process. If you are configuring multiarea IS-IS, specify a NET for each routing process. You can specify a name for a NET and for an address.
Step 5	is-type { level-1 level-1-2 level-2-only }	(Optional) You can configure the router to act as a Level 1 (station) router, a Level 2 (area) router for multi-area routing, or both (the default): level-1—act as a station router only level-1-2—act as both a station router and an area router level 2—act as an area router only

	Command	Purpose
Step 6	exit	Return to global configuration mode.
Step 7	interface <i>interface-id</i>	Specify an interface to route IS-IS, and enter interface configuration mode. If the interface is not already configured as a Layer 3 interface, enter the no switchport command to put it into Layer 3 mode.
Step 8	ip router isis [<i>area tag</i>]	Configure an IS-IS routing process for ISO CLNS on the interface and attach an area designator to the routing process.
Step 9	clns router isis [<i>area tag</i>]	Enable ISO CLNS on the interface.
Step 10	ip address <i>ip-address-mask</i>	Define the IP address for the interface. An IP address is required on all interfaces in an area enabled for IS-IS if any one interface is configured for IS-IS routing.
Step 11	end	Return to privileged EXEC mode.
Step 12	show isis [<i>area tag</i>] database detail	Verify your entries.
Step 13	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable IS-IS routing, use the **no router isis area-tag** router configuration command.

This example shows how to configure three routers to run conventional IS-IS as an IP routing protocol. In conventional IS-IS, all routers act as Level 1 and Level 2 routers (by default).

Router A

```
Switch(config)# clns routing
Switch(config)# router isis
Switch(config-router)# net 49.0001.0000.0000.000a.00
Switch(config-router)# exit
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# ip router isis
Switch(config-if)# clns router isis
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# ip router isis
Switch(config-if)# clns router isis
Switch(config-router)# exit
```

Router B

```
Switch(config)# clns routing
Switch(config)# router isis
Switch(config-router)# net 49.0001.0000.0000.000b.00
Switch(config-router)# exit
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# ip router isis
Switch(config-if)# clns router isis
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# ip router isis
Switch(config-if)# clns router isis
Switch(config-router)# exit
```

Router C

```
Switch(config)# clns routing
Switch(config)# router isis
Switch(config-router)# net 49.0001.0000.0000.000c.00
Switch(config-router)# exit
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# ip router isis
Switch(config-if)# clns router isis
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# ip router isis
Switch(config-if)# clns router isis
Switch(config-router)# exit
```

To configure multi-area routing, you configure Router B and Router C to act as Level 1 areas by entering the **is-type level-1** global configuration command.

Configuring IS-IS Global Parameters

These are some optional IS-IS global parameters that you can configure:

- You can force a default route into an IS-IS routing domain by configuring a default route controlled by a route map. You can also specify other filtering options configurable under a route map.
- You can configure the router to ignore IS-IS LSPs that are received with internal checksum errors or to purge corrupted LSPs, which causes the initiator of the LSP to regenerate it.
- You can assign passwords to areas and domains.
- You can create aggregate addresses that are represented in the routing table by a summary address (route-summarization). Routes learned from other routing protocols can also be summarized. The metric used to advertise the summary is the smallest metric of all the specific routes.
- You can set an overload bit.
- You can configure the LSP refresh interval and the maximum time that an LSP can remain in the router database without a refresh
- You can set the throttling timers for LSP generation, shortest path first computation, and partial route computation.
- You can configure the switch to generate a log message when an IS-IS adjacency changes state (up or down).
- If a link in the network has a maximum transmission unit (MTU) size of less than 1500 bytes, you can lower the LSP MTU so that routing will still occur.
- The partition avoidance router configuration command prevents an area from becoming partitioned when full connectivity is lost among a Level1-2 border router, adjacent Level 1 routers, and end hosts.

Beginning in privileged EXEC mode, follow these steps to configure IS-IS parameters:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	clns routing	Enable ISO connectionless routing on the switch.
Step 3	router isis	Specify the IS-IS routing protocol and enter router configuration mode.

	Command	Purpose
Step 4	default-information originate [route-map <i>map-name</i>]	(Optional) Force a default route into the IS-IS routing domain. If you enter route-map <i>map-name</i> , the routing process generates the default route if the route map is satisfied.
Step 5	ignore-lsp-errors	(Optional) Configure the router to ignore LSPs with internal checksum errors, instead of purging the LSPs. This command is enabled by default (corrupted LSPs are dropped). To purge the corrupted LSPs, enter the no ignore-lsp-errors router configuration command.
Step 6	area-password <i>password</i>	(Optional) Configure the area authentication password, which is inserted in Level 1 (station router level) LSPs.
Step 7	domain-password <i>password</i>	(Optional) Configure the routing domain authentication password, which is inserted in Level 2 (area router level) LSPs.
Step 8	summary-address <i>address mask</i> [level-1 level-1-2 level-2]	(Optional) Create a summary of addresses for a given level.
Step 9	set-overload-bit [on-startup { <i>seconds</i> wait-for-bgp }]	(Optional) Set an overload bit (a hippity bit) to allow other routers to ignore the router in their shortest path first (SPF) calculations if the router is having problems. • (Optional) on-startup—sets the overload bit only on startup. If on-startup is not specified, the overload bit is set immediately and remains set until you enter the no set-overload-bit command. If on-startup is specified, you must enter a number of seconds or wait-for-bgp. • <i>seconds</i>—When the on-startup keyword is configured, causes the overload bit to be set upon system startup and remain set for this number of seconds. The range is from 5 to 86400 seconds. • wait-for-bgp—When the on-startup keyword is configured, causes the overload bit to be set upon system startup and remain set until BGP has converged. If BGP does not signal IS-IS that it is converged, IS-IS will turn off the overload bit after 10 minutes.
Step 10	lsp-refresh-interval <i>seconds</i>	(Optional) Set an LSP refresh interval in seconds. The range is from 1 to 65535 seconds. The default is to send LSP refreshes every 900 seconds (15 minutes).
Step 11	max-lsp-lifetime <i>seconds</i>	(Optional) Set the maximum time that LSP packets remain in the router database without being refreshed. The range is from 1 to 65535 seconds. The default is 1200 seconds (20 minutes). After the specified time interval, the LSP packet is deleted.
Step 12	lsp-gen-interval [level-1 level-2] <i>lsp-max-wait</i> [<i>lsp-initial-wait</i> <i>lsp-second-wait</i>]	(Optional) Set the IS-IS LSP generation throttling timers: • <i>lsp-max-wait</i>—the maximum interval (in seconds) between two consecutive occurrences of an LSP being generated. The range is 1 to 120, the default is 5. • <i>lsp-initial-wait</i>—the initial LSP generation delay (in milliseconds). The range is 1 to 10000; the default is 50. • <i>lsp-second-wait</i>—the hold time between the first and second LSP generation (in milliseconds). The range is 1 to 10000; the default is 5000.

	Command	Purpose
Step 13	spf-interval [level-1 level-2] <i>spf-max-wait</i> [<i>spf-initial-wait</i> <i>spf-second-wait</i>]	(Optional) Sets IS-IS shortest path first (SPF) throttling timers. <i>spf-max-wait</i>—the maximum interval between consecutive SFPs (in seconds). The range is 1 to 120, the default is 10. <i>spf-initial-wait</i>—the initial SFP calculation after a topology change (in milliseconds). The range is 1 to 10000; the default is 5500. <i>spf-second-wait</i>—the holdtime between the first and second SFP calculation (in milliseconds). The range is 1 to 10000; the default is 5500.
Step 14	prc-interval <i>prc-max-wait</i> [<i>prc-initial-wait</i> <i>prc-second-wait</i>]	(Optional) Sets IS-IS partial route computation (PRC) throttling timers. <i>prc-max-wait</i>—the maximum interval (in seconds) between two consecutive PRC calculations. The range is 1 to 120; the default is 5. <i>prc-initial-wait</i>—the initial PRC calculation delay (in milliseconds) after a topology change. The range is 1 to 10,000; the default is 2000. <i>prc-second-wait</i>—the hold time between the first and second PRC calculation (in milliseconds). The range is 1 to 10,000; the default is 5000.
Step 15	log-adjacency-changes [all]	(Optional) Set the router to log IS-IS adjacency state changes. Enter all to include all changes generated by events that are not related to the Intermediate System-to-Intermediate System Hellos, including End System-to-Intermediate System PDUs and link state packets (LSPs).
Step 16	lsp-mtu <i>size</i>	(Optional) Specify the maximum LSP packet size in bytes. The range is 128 to 4352; the default is 1497 bytes. Note If any link in the network has a reduced MTU size, you must change the LSP MTU size on all routers in the network.
Step 17	partition avoidance	(Optional) Causes an IS-IS Level 1-2 border router to stop advertising the Level 1 area prefix into the Level 2 backbone when full connectivity is lost among the border router, all adjacent level 1 routers, and end hosts.
Step 18	end	Return to privileged EXEC mode.
Step 19	show clns	Verify your entries.
Step 20	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable default route generation, use the **no default-information originate** router configuration command. Use the **no area-password** or **no domain-password** router configuration command to disable passwords. To disable LSP MTU settings, use the **no lsp mtu** router configuration command. To return to the default conditions for summary addressing, LSP refresh interval, LSP lifetime, LSP timers, SFP timers, and PRC timers, use the **no** form of the commands. Use the **no partition avoidance** router configuration command to disable the output format.

Configuring IS-IS Interface Parameters

You can optionally configure certain interface-specific IS-IS parameters, independently from other attached routers. However, if you change some values from the defaults, such as multipliers and time intervals, it makes sense to also change them on multiple routers and interfaces. Most of the interface parameters can be configured for level 1, level 2, or both.

These are some interface level parameters you can configure:

- The default metric on the interface, which is used as a value for the IS-IS metric and assigned when there is no quality of service (QoS) routing performed.
- The hello interval (length of time between hello packets sent on the interface) or the default hello packet multiplier used on the interface to determine the hold time sent in IS-IS hello packets. The hold time determines how long a neighbor waits for another hello packet before declaring the neighbor down. This determines how quickly a failed link or neighbor is detected so that routes can be recalculated. Change the hello-multiplier in circumstances where hello packets are lost frequently and IS-IS adjacencies are failing unnecessarily. You can raise the hello multiplier and lower the hello interval correspondingly to make the hello protocol more reliable without increasing the time required to detect a link failure.
- Other time intervals:
 - Complete sequence number PDU (CSNP) interval. CSNPs are sent by the designated router to maintain database synchronization
 - Retransmission interval. This is the time between retransmission of IS-IS LSPs for point-to-point links.
 - IS-IS LSP retransmission throttle interval. This is the maximum rate (number of milliseconds between packets) at which IS-IS LSPs are re-sent on point-to-point links. This interval is different from the retransmission interval, which is the time between successive retransmissions of the *same* LSP
- Designated router election priority, which allows you to reduce the number of adjacencies required on a multiaccess network, which in turn reduces the amount of routing protocol traffic and the size of the topology database.
- The interface circuit type, which is the type of adjacency desired for neighbors on the specified interface
- Password authentication for the interface

Beginning in privileged EXEC mode, follow these steps to configure IS-IS interface parameters:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to be configured and enter interface configuration mode. If the interface is not already configured as a Layer 3 interface, enter the no switchport command to put it into Layer 3 mode.
Step 3	isis metric <i>default-metric</i> [level-1 level-2]	(Optional) Configure the metric (or cost) for the specified interface. The range is from 0 to 63. The default is 10. If no level is entered, the default is to apply to both Level 1 and Level 2 routers.
Step 4	isis hello-interval { <i>seconds</i> minimal } [level-1 level-2]	(Optional) Specify the length of time between hello packets sent by the switch. By default, a value three times the hello interval <i>seconds</i> is advertised as the <i>holdtime</i> in the hello packets sent. With smaller hello intervals, topological changes are detected faster, but there is more routing traffic. • minimal—causes the system to compute the hello interval based on the hello multiplier so that the resulting hold time is 1 second. • <i>seconds</i>—the range is from 1 to 65535. The default is 10 seconds.

	Command	Purpose
Step 5	isis hello-multiplier <i>multiplier</i> [level-1 level-2]	(Optional) Specify the number of IS-IS hello packets a neighbor must miss before the router should declare the adjacency as down. The range is from 3 to 1000. The default is 3. Using a smaller hello-multiplier causes fast convergence, but can result in more routing instability.
Step 6	isis csnp-interval <i>seconds</i> [level-1 level-2]	(Optional) Configure the IS-IS complete sequence number PDU (CSNP) interval for the interface. The range is from 0 to 65535. The default is 10 seconds.
Step 7	isis retransmit-interval <i>seconds</i>	(Optional) Configure the number of seconds between retransmission of IS-IS LSPs for point-to-point links. The value you specify should be an integer greater than the expected round-trip delay between any two routers on the network. The range is from 0 to 65535. The default is 5 seconds.
Step 8	isis retransmit-throttle-interval <i>milliseconds</i>	(Optional) Configure the IS-IS LSP retransmission throttle interval, which is the maximum rate (number of milliseconds between packets) at which IS-IS LSPs will be re-sent on point-to-point links. The range is from 0 to 65535. The default is determined by the isis lsp-interval command.
Step 9	isis priority <i>value</i> [level-1 level-2]	(Optional) Configure the priority to use for designated router election. The range is from 0 to 127. The default is 64.
Step 10	isis circuit-type { level-1 level-1-2 level-2-only }	(Optional) Configure the type of adjacency desired for neighbors on the specified interface (specify the interface circuit type). • level-1—a Level 1 adjacency is established if there is at least one area address common to both this node and its neighbors. • level-1-2—a Level 1 and 2 adjacency is established if the neighbor is also configured as both Level 1 and Level 2 and there is at least one area in common. If there is no area in common, a Level 2 adjacency is established. This is the default. • level 2—a Level 2 adjacency is established. If the neighbor router is a Level 1 router, no adjacency is established.
Step 11	isis password <i>password</i> [level-1 level-2]	(Optional) Configure the authentication password for an interface. By default, authentication is disabled. Specifying Level 1 or Level 2 enables the password only for Level 1 or Level 2 routing, respectively. If you do not specify a level, the default is Level 1 and Level 2.
Step 12	end	Return to privileged EXEC mode.
Step 13	show clns interface <i>interface-id</i>	Verify your entries.
Step 14	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default settings, use the **no** forms of the commands.

Monitoring and Maintaining ISO IGRP and IS-IS

You can remove all contents of a CLNS cache or remove information for a particular neighbor or route. You can display specific CLNS or IS-IS statistics, such as the contents of routing tables, caches, and databases. You can also display information about specific interfaces, filters, or neighbors.

Table 28-15 lists the privileged EXEC commands for clearing and displaying ISO CLNS and IS-IS routing. For explanations of the display fields, refer to the *Cisco IOS Apollo Domain, Banyan VINES, DECnet, ISO CLNS and XNS Command Reference for Release 12.1*, use the Cisco IOS command reference master index, or search online.

Table 28-15 ISO CLNS and IS-IS Clear and Show Commands

Command	Purpose
clear clns cache	Clear and reinitialize the CLNS routing cache.
clear clns es-neighbors	Remove end system (ES) neighbor information from the adjacency database.
clear clns is-neighbors	Remove intermediate system (IS) neighbor information from the adjacency database.
clear clns neighbors	Remove CLNS neighbor information from the adjacency database.
clear clns route	Remove dynamically derived CLNS routing information.
show clns	Display information about the CLNS network.
show clns cache	Display the entries in the CLNS routing cache.
show clns es-neighbors	Display ES neighbor entries, including the associated areas.
show clns filter-expr	Display filter expressions.
show clns filter-set	Display filter sets.
show clns interface <i>[interface-id]</i>	Display the CLNS-specific or ES-IS information about each interface.
show clns neighbor	Display information about IS-IS neighbors.
show clns protocol	List the protocol-specific information for each IS-IS or ISO IGRP routing process in this router.
show clns route	Display all the destinations to which this router knows how to route CLNS packets.
show clns traffic	Display information about the CLNS packets this router has seen.
show ip route isis	Display the current state of the ISIS IP routing table.
show isis database	Display the IS-IS link-state database.
show isis routes	Display the IS-IS Level 1 routing table.
show isis spf-log	Display a history of the shortest path first (SPF) calculations for IS-IS.
show isis topology	Display a list of all connected routers in all areas.
show route-map	Display all route maps configured or only the one specified.
trace clns <i>destination</i>	Discover the paths taken to a specified destination by packets in the network.
which-route <i>{nsap-address clns-name}</i>	Display the routing table in which the specified CLNS destination is found.

Configuring Multi-VRF CE

Virtual Private Networks (VPNs) provide a secure way for customers to share bandwidth over a service-provider backbone network. A VPN is a collection of sites sharing a common routing table. A customer site is connected to the service-provider network by one or more interfaces, and the service provider associates each interface with a VPN routing table, called a VPN routing/forwarding (VRF) table. In a multiprotocol label switching (MPLS) VPN deployment, all customers have their own customer-edge (CE) device, and every CE has a dedicated line connected to a provider-edge (PE) router.

With multiple VPN routing/forwarding (multi-VRF), multiple customers can share one CE, and only one physical link is needed between the CE and the PE. This is accomplished by extending some PE functionality to the shared CE devices. The shared CE maintains separate VRF tables for each customer and switches or routes packets for each customer based on the CE routing table. Multi-VRF CE allows a service provider to support two or more VPNs with overlapping IP addresses.

This section includes these topics:

- [Understanding Multi-VRF CE, page 28-76](#)
- [Default Multi-VRF CE Configuration, page 28-78](#)
- [Multi-VRF CE Configuration Guidelines, page 28-79](#)
- [Configuring VRFs, page 28-80](#)
- [Configuring a VPN Routing Session, page 28-80](#)
- [Configuring BGP PE to CE Routing Sessions, page 28-81](#)
- [Multi-VRF CE Configuration Example, page 28-82](#)
- [Displaying Multi-VRF CE Status, page 28-86](#)

Understanding Multi-VRF CE

Multi-VRF CE is a feature that allows a service provider to support two or more VPNs, where IP addresses can be overlapped among the VPNs. Multi-VRF CE uses input interfaces to distinguish routes for different VPNs and forms virtual packet-forwarding tables by associating one or more Layer 3 interfaces with each VRF. Interfaces in a VRF can be either physical, such as Ethernet ports, or logical, such as VLAN SVIs, but an interface cannot belong to more than one VRF at any time.

**Note**

Multi-VRF CE interfaces must be Layer 3 interfaces.

Multi-VRF CE includes these devices:


- CE devices provide customers access to the service-provider network over a data link to one or more PE routers. The CE device advertises the site's local routes to the router and learns the remote VPN routes from it. A Catalyst 3750 Metro switch can be a CE.
- PE routers exchange routing information with CE devices by using static routing or a routing protocol such as BGP, RIPv2, OSPF, IS-IS, or EIGRP. The PE is only required to maintain VPN routes for those VPNs to which it is directly attached, eliminating the need for the PE to maintain all of the service-provider VPN routes. Each PE router maintains a VRF for each of its directly connected sites. Multiple interfaces on a PE router can be associated with a single VRF if all of these sites participate in the same VPN. Each VPN is mapped to a specified VRF. After learning local VPN routes from CEs, a PE router exchanges VPN routing information with other PE routers by using internal BGP (IBPG). A Catalyst 3750 Metro switch would typically be used as a PE.

- Provider routers or core routers are any routers in the service-provider network that do not attach to CE devices.

With multi-VRF CE, multiple customers can share one CE, and only one physical link is used between the CE and the PE. The shared CE maintains separate VRF tables for each customer and switches or routes packets for each customer based on its own routing table. Multi-VRF CE extends limited PE functionality to a CE device, giving it the ability to maintain separate VRF tables to extend the privacy and security of a VPN to the branch office.

Figure 28-6 shows a configuration using Catalyst 3750 Metro switches as multiple virtual CEs. This scenario is suited for customers who have low bandwidth requirements for their VPN service, for example, small companies. In this case, multi-VRF CE support is required in the Catalyst 3750 Metro switches. Because multi-VRF CE is a Layer 3 feature, each interface in a VRF must be a Layer 3 interface.

Figure 28-6 Catalyst 3750 Metro Switches Acting as Multiple Virtual CEs



When the CE switch receives a command to add a Layer 3 interface to a VRF, it sets up the appropriate mapping between the VLAN ID and the policy label in multi-VRF-CE-related data structures and adds the VLAN ID and policy label to the VLAN database.

When the multi-VRF CE is configured, the Layer 3 forwarding table is conceptually partitioned into two sections:

- The multi-VRF CE routing section contains the routes from different VPNs.
- The global routing section contains routes to non-VPN networks, such as the Internet.

VLAN IDs from different VRFs are mapped into different policy labels, which are used to distinguish the VRFs during processing. If no route is found in the multi-VRF CE section of the Layer 3 forwarding table, the global routing section is used to determine the forwarding path. For each new VPN route learned, the Layer 3 setup function retrieves the policy label by using the VLAN ID of the ingress port and inserts the policy label and new route to the multi-VRF CE routing section. If the packet is received from a routed port, the port internal VLAN ID number is used; if the packet is received from an SVI, the VLAN number is used.

This is the packet-forwarding process in a multi-VRF-CE-enabled network:

- When the switch receives a packet from a VPN, the switch looks up the routing table, based on the input policy label number. When a route is found, the switch forwards the packet to the PE.
- When the ingress PE receives a packet from the CE, it performs a VRF lookup. When a route is found, the router adds a corresponding MPLS label to the packet and sends it to the MPLS network.
- When an egress PE receives a packet from the network, it removes the label and uses the label to identify the correct VPN routing table. Then it performs the normal route lookup. When a route is found, it forwards the packet to the correct adjacency.
- When a CE receives a packet from an egress PE, it uses the input policy label to look up the correct VPN routing table. When a route is found, it forwards the packet within the VPN.

To configure VRF, you create a VRF table and specify the Layer 3 interface associated with the VRF. Then configure the routing protocols in the VPN and between the CE and the PE. BGP is the preferred routing protocol used to distribute VPN routing information across the provider's backbone.

The multi-VRF CE network has three major components:

- VPN route target communities—lists of all other members of a VPN community. You need to configure VPN route targets for each VPN community member.
- Multiprotocol BGP peering of VPN community PE routers—propagates VRF reachability information to all members of a VPN community. You need to configure BGP peering in all PE routers within a VPN community.
- VPN forwarding—transports all traffic between all VPN community members across a VPN service-provider network.

Default Multi-VRF CE Configuration

Table 28-16 shows the default VRF configuration.

Table 28-16 Default VRF Configuration

Feature	Default Setting
VRF	Disabled. No VRFs are defined.
Maps	No import maps, export maps, or route maps are defined.
VRF maximum routes	8000 (total number of routes supported in hardware).
Forwarding table	The default for an interface is the global routing table.

Multi-VRF CE Configuration Guidelines

These are considerations when configuring VRF in your network:

- A switch with multi-VRF CE is shared by multiple customers, and each customer has its own routing table.
- Because customers use different VRF tables, the same IP addresses can be reused. Overlapped IP addresses are allowed in different VPNs.
- Multi-VRF CE lets multiple customers share the same physical link between the PE and the CE. Trunk ports with multiple VLANs separate packets among customers. Each customer has its own VLAN.
- Multi-VRF CE does not support all MPLS-VRF functionality. It does not support label exchange, LDP adjacency, or labeled packets.
- For the PE router, there is no difference between using multi-VRF CE or using multiple CEs. In [Figure 28-6](#), multiple virtual Layer 3 interfaces are connected to the multi-VRF CE device.
- The switch supports configuring VRF by using physical ports, VLAN SVIs, or a combination of both. The SVIs can be connected through an access port or a trunk port.
- The switch supports a total of 26 VRFs and VPNs.
- A customer can use multiple VLANs as long as they do not overlap with those of other customers. A customer's VLANs are mapped to a specific routing table ID that is used to identify the appropriate routing tables stored on the switch.
- A switch using VRF can support one global network and up to seven VRFs.
- Most routing protocols (BGP, OSPF, IS-IS, RIP, and static routing) can be used between the CE and the PE. However, we recommend using external BGP (EBGP) for these reasons:
 - BGP does not require multiple algorithms to communicate with multiple CEs.
 - BGP is designed for passing routing information between systems run by different administrations.
 - BGP makes it easy to pass attributes of the routes to the CE.
- Multi-VRF-CE does not support IGRP and EIGRP.
- Multi-VRF CE does not affect the packet switching rate.
- VPN multicast is not supported.
- When multi-VRF CE is configured, you cannot assign the same Hot Standby Routing Protocol (HSRP) standby address to two different VPNs.
- VRFs are not compatible with the PBR template. If you configure the PBR template by entering the **sdm prefer routing-pbr** command, any preconfigured VRFs are removed from the configuration. PBR and VRFs cannot function on the same switch.

Configuring VRFs

Beginning in privileged EXEC mode, follow these steps to configure one or more VRFs. For complete syntax and usage information for the commands, refer to the switch command reference for this release and the *Cisco IOS Switching Services Command Reference for Release 12.1*.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip routing	Enable IP routing.
Step 3	ip vrf vrf-name	Name the VRF, and enter VRF configuration mode.
Step 4	rd route-distinguisher	Create a VRF table by specifying a route distinguisher. Enter either an AS number and an arbitrary number (xxx:y) or an IP address and arbitrary number (A.B.C.D:y)
Step 5	route-target {export import both} <i>route-target-ext-community</i>	Create a list of import, export, or import and export route target communities for the specified VRF. Enter either an AS system number and an arbitrary number (xxx:y) or an IP address and an arbitrary number (A.B.C.D:y). The <i>route-target-ext-community</i> should be the same as the <i>route-distinguisher</i> entered in Step 4.
Step 6	import map route-map	(Optional) Associate a route map with the VRF.
Step 7	interface interface-id	Enter interface configuration mode and specify the Layer 3 interface to be associated with the VRF. The interface can be a routed port or SVI.
Step 8	ip vrf forwarding vrf-name	Associate the VRF with the Layer 3 interface.
Step 9	end	Return to privileged EXEC mode.
Step 10	show ip vrf [brief detail interfaces] <i>[vrf-name]</i>	Verify the configuration. Display information about the configured VRFs.
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip vrf vrf-name** global configuration command to delete a VRF and to remove all interfaces from it. Use the **no ip vrf forwarding** interface configuration command to remove an interface from the VRF.

Configuring a VPN Routing Session

Routing within the VPN can be configured with any supported routing protocol (RIP, OSPF, IGRP, EIGRP, IS-IS or BGP) or with static routing. The configuration shown here is for OSPF, but the process is the same for other protocols.

Beginning in privileged EXEC mode, follow these steps to configure OSPF in the VPN:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router ospf process-id vrf vrf-name	Enable OSPF routing, specify a VPN forwarding table, and enter router configuration mode.

	Command	Purpose
Step 3	log-adjacency-changes	(Optional) Log changes in the adjacency state. This is the default state.
Step 4	redistribute bgp <i>autonomous-system-number</i> subnets	Set the switch to redistribute information from the BGP network to the OSPF network.
Step 5	network <i>network-number</i> area <i>area-id</i>	Define a network address and mask on which OSPF runs and the area ID for that network address.
Step 6	end	Return to privileged EXEC mode.
Step 7	show ip ospf <i>process-id</i>	Verify the configuration of the OSPF network.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no router ospf** *process-id* **vrf** *vrf-name* global configuration command to disassociate the VPN forwarding table from the OSPF routing process.

Configuring BGP PE to CE Routing Sessions

Beginning in privileged EXEC mode, follow these steps to configure a BGP PE to CE routing session:


	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system-number</i>	Configure the BGP routing process with the AS number passed to other BGP routers, and enter router configuration mode.
Step 3	network <i>network-number</i> mask <i>network-mask</i>	Specify a network and mask to announce using BGP.
Step 4	redistribute ospf <i>process-id</i> match internal	Set the switch to redistribute OSPF internal routes.
Step 5	network <i>network-number</i> area <i>area-id</i>	Define a network address and mask on which OSPF runs and the area ID for that network address.
Step 6	address-family ipv4 vrf <i>vrf-name</i>	Define BGP parameters for PE to CE routing sessions, and enter VRF address-family mode.
Step 7	neighbor <i>address</i> remote-as <i>as-number</i>	Define a BGP session between PE and CE routers.
Step 8	neighbor <i>address</i> activate	Activate the advertisement of the IPv4 address family.
Step 9	end	Return to privileged EXEC mode.
Step 10	show ip bgp [ipv4] [neighbors]	Verify BGP configuration.
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no router bgp** *autonomous-system-number* global configuration command to delete the BGP routing process. Use the command with keywords to delete routing characteristics.

Multi-VRF CE Configuration Example

Figure 28-7 is a simplified example of the physical connections in a network similar to that in Figure 28-6. OSPF is the protocol used in VPN1, VPN2, and the global network. BGP is used in the CE-to-PE connections. The examples following the illustration show how to configure a Catalyst 3750 Metro switch as CE Switch A and the VRF configuration for customer switches D and F. Commands for configuring CE Switch C and the other customer switches are not included, but would be similar to those shown. The example also includes commands for configuring traffic to Switch A for a Catalyst 6000 or Catalyst 6500 switch acting as PE router.

Figure 28-7 Multi-VRF CE Configuration Example



Configuring Switch A

On Switch A, enable routing, and configure VRF.

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip routing
Switch(config)# ip vrf v11
Switch(config-vrf)# rd 800:1
Switch(config-vrf)# route-target export 800:1
Switch(config-vrf)# route-target import 800:1
Switch(config-vrf)# exit
Switch(config)# ip vrf v12
Switch(config-vrf)# rd 800:2
Switch(config-vrf)# route-target export 800:2
Switch(config-vrf)# route-target import 800:2
Switch(config-vrf)# exit
```


Configure the loopback and physical interfaces on Switch A. Gigabit Ethernet port 1 is a trunk connection to the PE. Fast Ethernet ports 8 and 11 connect to VPNs:

```
Switch(config)# interface loopback1
Switch(config-if)# ip vrf forwarding v11
Switch(config-if)# ip address 8.8.1.8 255.255.255.0
Switch(config-if)# exit

Switch(config)# interface loopback2
Switch(config-if)# ip vrf forwarding v12
Switch(config-if)# ip address 8.8.2.8 255.255.255.0
Switch(config-if)# exit

Switch(config)# interface gigabitethernet1/1/1
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport mode trunk
Switch(config-if)# no ip address
Switch(config-if)# exit

Switch(config)# interface fastethernet1/0/8
Switch(config-if)# switchport access vlan 208
Switch(config-if)# no ip address
Switch(config-if)# exit

Switch(config)# interface fastethernet1/0/11
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport mode trunk
Switch(config-if)# no ip address
Switch(config-if)# exit
```

Configure the VLANs used on Switch A. VLAN 10 is used by VRF 11 between the CE and the PE. VLAN 20 is used by VRF 12 between the CE and the PE. VLANs 118 and 208 are used for VRF for the VPNs that include Switch F and Switch D, respectively:

```
Switch(config)# interface vlan10
Switch(config-if)# ip vrf forwarding v11
Switch(config-if)# ip address 38.0.0.8 255.255.255.0
Switch(config-if)# exit

Switch(config)# interface vlan20
Switch(config-if)# ip vrf forwarding v12
Switch(config-if)# ip address 83.0.0.8 255.255.255.0
Switch(config-if)# exit

Switch(config)# interface vlan118
Switch(config-if)# ip vrf forwarding v12
Switch(config-if)# ip address 118.0.0.8 255.255.255.0
Switch(config-if)# exit

Switch(config)# interface vlan208
Switch(config-if)# ip vrf forwarding v11
Switch(config-if)# ip address 208.0.0.8 255.255.255.0
Switch(config-if)# exit
```

Configure OSPF routing in VPN1 and VPN2.

```
Switch(config)# router ospf 1 vrf v11
Switch(config-router)# redistribute bgp 800 subnets
Switch(config-router)# network 208.0.0.0 0.0.0.255 area 0
Switch(config-router)# exit
Switch(config)# router ospf 2 vrf v12
Switch(config-router)# redistribute bgp 800 subnets
Switch(config-router)# network 118.0.0.0 0.0.0.255 area 0
Switch(config-router)# exit
```

Configure BGP for CE to PE routing.

```
Switch(config)# router bgp 800
Switch(config-router)# address-family ipv4 vrf vl2
Switch(config-router-af)# redistribute ospf 2 match internal
Switch(config-router-af)# neighbor 83.0.0.3 remote-as 100
Switch(config-router-af)# neighbor 83.0.0.3 activate
Switch(config-router-af)# network 8.8.2.0 mask 255.255.255.0
Switch(config-router-af)# exit
```

```
Switch(config-router)# address-family ipv4 vrf vl1
Switch(config-router-af)# redistribute ospf 1 match internal
Switch(config-router-af)# neighbor 38.0.0.3 remote-as 100
Switch(config-router-af)# neighbor 38.0.0.3 activate
Switch(config-router-af)# network 8.8.1.0 mask 255.255.255.0
Switch(config-router-af)# end
```

Configuring Switch D

Switch D belongs to VPN 1 and is connected to Switch A by using these commands.

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip routing
Switch(config)# interface fastethernet1/0/2
Switch(config-if)# no switchport
Switch(config-if)# ip address 208.0.0.20 255.255.255.0
Switch(config-if)# exit
```

```
Switch(config)# router ospf 101
Switch(config-router)# network 208.0.0.0 0.0.0.255 area 0
Switch(config-router)# end
```

Configuring Switch F

Switch F belongs to VPN 2 and is connected to Switch A by using these commands.

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip routing
Switch(config)# interface fastethernet1/0/1
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport mode trunk
Switch(config-if)# no ip address
Switch(config-if)# exit
```

```
Switch(config)# interface Vlan118
Switch(config-if)# ip address 118.0.0.11 255.255.255.0
Switch(config-if)# exit
```

```
Switch(config)# router ospf 101
Switch(config-router)# network 118.0.0.0 0.0.0.255 area 0
Switch(config-router)# end
```

Configuring the PE Switch B

On Switch B (the PE router), these commands configure the connections to the CE device, Switch A.

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip vrf v1
Router(config-vrf)# rd 100:1
Router(config-vrf)# route-target export 100:1
Router(config-vrf)# route-target import 100:1
Router(config-vrf)# exit

Router(config)# ip vrf v2
Router(config-vrf)# rd 100:2
Router(config-vrf)# route-target export 100:2
Router(config-vrf)# route-target import 100:2
Router(config-vrf)# exit

Router(config)# ip cef
Router(config)# interface Loopback1
Router(config-if)# ip vrf forwarding v1
Router(config-if)# ip address 3.3.1.3 255.255.255.0
Router(config-if)# exit

Router(config)# interface Loopback2
Router(config-if)# ip vrf forwarding v2
Router(config-if)# ip address 3.3.2.3 255.255.255.0
Router(config-if)# exit

Router(config)# interface gigabitethernet1/0.10
Router(config-if)# encapsulation dot1q 10
Router(config-if)# ip vrf forwarding v1
Router(config-if)# ip address 38.0.0.3 255.255.255.0
Router(config-if)# exit

Router(config)# interface gigabitethernet1/0.20
Router(config-if)# encapsulation dot1q 20
Router(config-if)# ip vrf forwarding v2
Router(config-if)# ip address 83.0.0.3 255.255.255.0
Router(config-if)# exit

Router(config)# router bgp 100
Router(config-router)# address-family ipv4 vrf v2
Router(config-router-af)# neighbor 83.0.0.8 remote-as 800
Router(config-router-af)# neighbor 83.0.0.8 activate
Router(config-router-af)# network 3.3.2.0 mask 255.255.255.0
Router(config-router-af)# exit
Router(config-router)# address-family ipv4 vrf v1
Router(config-router-af)# neighbor 83.0.0.8 remote-as 800
Router(config-router-af)# neighbor 83.0.0.8 activate
Router(config-router-af)# network 3.3.1.0 mask 255.255.255.0
Router(config-router-af)# end
```

Displaying Multi-VRF CE Status

You can use the privileged EXEC commands in [Table 28-7](#) to display information about multi-VRF CE configuration and status.

Table 28-17 Show IP OSPF Statistics Commands

Command	Purpose
<code>show ip protocols vrf vrf-name</code>	Display routing protocol information associated with a VRF.
<code>show ip route vrf vrf-name [connected] [protocol [as-number]] [list] [mobile] [odr] [profile] [static] [summary] [supernets-only]</code>	Display IP routing table information associated with a VRF.
<code>show ip vrf [brief detail interfaces] [vrf-name]</code>	Display information about the defined VRF instances.

For more information about the information in the displays, refer to the *Cisco IOS Switching Services Command Reference for Release 12.1*.

Configuring Protocol-Independent Features

This section describes how to configure IP routing protocol-independent features. For a complete description of the IP routing protocol-independent commands in this chapter, refer to the “IP Routing Protocol-Independent Commands” chapter of the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

This section includes these procedures:

- [Configuring Cisco Express Forwarding, page 28-86](#)
- [Configuring the Number of Equal-Cost Routing Paths, page 28-87](#)
- [Configuring Static Unicast Routes, page 28-88](#)
- [Specifying Default Routes and Networks, page 28-89](#)
- [Using Route Maps to Redistribute Routing Information, page 28-90](#)
- [Configuring Policy-Based Routing, page 28-94](#)
- [Filtering Routing Information, page 28-96](#)
- [Managing Authentication Keys, page 28-99](#)

Configuring Cisco Express Forwarding

Cisco Express Forwarding (CEF) is a Layer 3 IP switching technology used to optimize network performance. CEF implements an advanced IP look-up and forwarding algorithm to deliver maximum Layer 3 switching performance. CEF is less CPU-intensive than fast switching route caching, allowing more CPU processing power to be dedicated to packet forwarding. In dynamic networks, fast switching cache entries are frequently invalidated because of routing changes, which can cause traffic to be process switched using the routing table, instead of fast switched using the route cache. CEF uses the Forwarding Information Base (FIB) lookup table to perform destination-based switching of IP packets.

The two main components in CEF are the distributed FIB and the distributed adjacency tables.

- The FIB is similar to a routing table or information base and maintains a mirror image of the forwarding information in the IP routing table. When routing or topology changes occur in the network, the IP routing table is updated, and those changes are reflected in the FIB. The FIB maintains next-hop address information based on the information in the IP routing table. Because the FIB contains all known routes that exist in the routing table, CEF eliminates route cache maintenance, is more efficient for switching traffic, and is not affected by traffic patterns.
- Nodes in the network are said to be adjacent if they can reach each other with a single hop across a link layer. CEF uses adjacency tables to prepend Layer 2 addressing information. The adjacency table maintains Layer 2 next-hop addresses for all FIB entries.

CEF is enabled globally by default. If for some reason it is disabled, you can re-enable it by using the **ip cef** global configuration command.

The default configuration is CEF enabled on all Layer 3 interfaces.



Caution

Although the **no ip route-cache cef** interface configuration command to disable CEF on an interface is visible in the CLI, we strongly recommend that you do not disable CEF on interfaces.

Beginning in privileged EXEC mode, follow these steps to enable CEF globally and on an interface in case, if, for some reason, it has been disabled:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip cef	Enable CEF operation.
Step 3	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 4	ip route-cache cef	Enable CEF on the interface.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip cef	Display the CEF status on all interfaces.
Step 7	show cef interface [<i>interface-id</i>]	Display detailed CEF information for all interfaces or the specified interface.
Step 8	show adjacency	Display CEF adjacency table information.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Configuring the Number of Equal-Cost Routing Paths

When a router has two or more routes to the same network with the same metrics, these routes can be thought of as having an equal cost. The term *parallel path* is another way to refer to occurrences of equal-cost routes in a routing table. If a router has two or more equal-cost paths to a network, it can use them concurrently. Parallel paths provide redundancy in case of a circuit failure and also enable a router to load balance packets over the available paths for more efficient use of available bandwidth.

Although the router automatically learns about and configures equal-cost routes, you can control the maximum number of parallel paths supported by an IP routing protocol in its routing table.

Beginning in privileged EXEC mode, follow these steps to change the maximum number of parallel paths installed in a routing table from the default:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router {bgp eigrp igmp isis iso-igrp ospf rip}	Enter router configuration mode.
Step 3	maximum-paths <i>maximum</i>	Set the maximum number of parallel paths for the protocol routing table. The range is from 1 to 8; the default is 4 for most IP routing protocols, but only 1 for BGP.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip protocols	Verify the setting in the <i>Maximum path</i> field.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no maximum-paths** router configuration command to restore the default value.

Configuring Static Unicast Routes

Static unicast routes are user-defined routes that cause packets moving between a source and a destination to take a specified path. Static routes can be important if the router cannot build a route to a particular destination and are useful for specifying a gateway of last resort to which all unroutable packets are sent.

Beginning in privileged EXEC mode, follow these steps to configure a static route:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip route <i>prefix mask {address interface} [distance]</i>	Establish a static route.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip route	Display the current state of the routing table to verify the configuration.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip route *prefix mask {address | interface}*** global configuration command to remove a static route.

The switch retains static routes until you remove them. However, you can override static routes with dynamic routing information by assigning administrative distance values. Each dynamic routing protocol has a default administrative distance, as listed in [Table 28-18](#). If you want a static route to be overridden by information from a dynamic routing protocol, set the administrative distance of the static route higher than that of the dynamic protocol.

Table 28-18 Dynamic Routing Protocol Default Administrative Distances

Route Source	Default Distance
Connected interface	0
Static route	1
Enhanced IRGP summary route	5
External BGP	20
Internal Enhanced IGRP	90
IGRP	100
OSPF	110
RIP	120
Internal BGP	200
Unknown	225

Static routes that point to an interface are advertised through RIP, IGRP, and other dynamic routing protocols, whether or not static **redistribute** router configuration commands were specified for those routing protocols. These static routes are advertised because static routes that point to an interface are considered in the routing table to be connected and hence lose their static nature. However, if you define a static route to an interface that is not one of the networks defined in a network command, no dynamic routing protocols advertise the route unless a **redistribute** static command is specified for these protocols.

When an interface goes down, all static routes through that interface are removed from the IP routing table. When the software can no longer find a valid next hop for the address specified as the forwarding router's address in a static route, the static route is also removed from the IP routing table.

Specifying Default Routes and Networks

A router might not be able to determine the routes to all other networks. To provide complete routing capability, you can use some routers as smart routers and give the remaining routers default routes to the smart router. (Smart routers have routing table information for the entire internetwork.) These default routes can be dynamically learned or can be configured in the individual routers. Most dynamic interior routing protocols include a mechanism for causing a smart router to generate dynamic default information that is then forwarded to other routers.

If a router has a directly connected interface to the specified default network, the dynamic routing protocols running on that device generate a default route. In RIP, it advertises the pseudonetwork 0.0.0.0. In IGRP, the network itself is advertised and flagged as an exterior route.

A router that is generating the default for a network also might need a default of its own. One way a router can generate its own default is to specify a static route to the network 0.0.0.0 through the appropriate device.

Beginning in privileged EXEC mode, follow these steps to define a static route to a network as the static default route:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip default-network <i>network number</i>	Specify a default network.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip route	Display the selected default route in the gateway of last resort display.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip default-network** *network number* global configuration command to remove the route.

When default information is passed through a dynamic routing protocol, no further configuration is required. The system periodically scans its routing table to choose the optimal default network as its default route. In IGRP networks, there might be several candidate networks for the system default. Cisco routers use administrative distance and metric information to determine the default route or the gateway of last resort.

If dynamic default information is not being passed to the system, candidates for the default route are specified with the **ip default-network** global configuration command. If this network appears in the routing table from any source, it is flagged as a possible choice for the default route. If the router has no interface on the default network, but does have a path to it, the network is considered as a possible candidate, and the gateway to the best default path becomes the gateway of last resort.

Using Route Maps to Redistribute Routing Information

The switch can run multiple routing protocols simultaneously, and it can redistribute information from one routing protocol to another. For example, you can instruct the switch to readvertise IGRP-derived routes by using RIP or to readvertise static routes by using IGRP. Redistributing information from one routing protocol to another applies to all supported IP-based routing protocols.

You can also conditionally control the redistribution of routes between routing domains by defining enhanced packet filters or route maps between the two domains. The **match** and **set** route-map configuration commands define the condition portion of a route map. The **match** command specifies that a criterion must be matched; the **set** command specifies an action to be taken if the routing update meets the conditions defined by the match command. Although redistribution is a protocol-independent feature, some of the **match** and **set** route-map configuration commands are specific to a particular protocol.

One or more **match** commands and one or more **set** commands follow a **route-map** command. If there are no **match** commands, everything matches. If there are no **set** commands, nothing is done, other than the match. Therefore, you need at least one **match** or **set** command.

You can also identify route-map statements as **permit** or **deny**. If the statement is marked as a deny, the packets meeting the match criteria are sent back through the normal forwarding channels (destination-based routing). If the statement is marked as permit, set clauses are applied to packets meeting the match criteria. Packets that do not meet the match criteria are forwarded through the normal routing channel.

**Note**

Although each of Steps 3 through 14 in the following section is optional, you must enter at least one **match** route-map configuration command and one **set** route-map configuration command.

Beginning in privileged EXEC mode, follow these steps to configure a route map for redistribution:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	route-map <i>map-tag</i> [permit deny] [<i>sequence number</i>]	Define any route maps used to control redistribution and enter route-map configuration mode. <i>map-tag</i> —A meaningful name for the route map. The redistribute router configuration command uses this name to reference this route map. Multiple route maps might share the same map tag name. (Optional) If permit is specified and the match criteria are met for this route map, the route is redistributed as controlled by the set actions. If deny is specified, the route is not redistributed. <i>sequence number</i> (Optional)— Number that indicates the position a new route map is to have in the list of route maps already configured with the same name.
Step 3	match as-path <i>path-list-number</i>	Match a BGP AS path access list.
Step 4	match community-list <i>community-list-number</i> [exact]	Match a BGP community list.
Step 5	match ip address { <i>access-list-number</i> <i>access-list-name</i> } [... <i>access-list-number</i> ... <i>access-list-name</i>]	Match a standard access list by specifying the name or number. It can be an integer from 1 to 199.
Step 6	match metric <i>metric-value</i>	Match the specified route metric. The <i>metric-value</i> can be an IGRP five-part metric with a specified value from 0 to 4294967295.
Step 7	match ip next-hop { <i>access-list-number</i> <i>access-list-name</i> } [... <i>access-list-number</i> ... <i>access-list-name</i>]	Match a next-hop router address passed by one of the access lists specified (numbered from 1 to 199).
Step 8	match tag <i>tag value</i> [... <i>tag-value</i>]	Match the specified tag value in a list of one or more route tag values. Each can be an integer from 0 to 4294967295.
Step 9	match interface <i>type number</i> [... <i>type number</i>]	Match the specified next hop route out one of the specified interfaces.
Step 10	match ip route-source { <i>access-list-number</i> <i>access-list-name</i> } [... <i>access-list-number</i> ... <i>access-list-name</i>]	Match the address specified by the specified advertised access lists.
Step 11	match route-type { local internal external [type-1 type-2]}	Match the specified route-type : • local—Locally generated BGP routes. • internal—OSPF intra-area and interarea routes or EIGRP internal routes. • external—OSPF external routes (Type 1 or Type 2) or EIGRP external routes.

	Command	Purpose
Step 12	set dampening <i>half-life reuse suppress max-suppress-time</i>	Set BGP route dampening factors.
Step 13	set local-preference <i>value</i>	Assign a value to a local BGP path.
Step 14	set origin { igp egp <i>as</i> incomplete }	Set the BGP origin code.
Step 15	set as-path { tag prepend <i>as-path-string</i> }	Modify the BGP autonomous system path.
Step 16	set level { level-1 level-2 level-1-2 stub-area backbone }	Set the level for routes that are advertised into the specified area of the routing domain. The stub-area and backbone are OSPF NSSA and backbone areas.
Step 17	set metric <i>metric value</i>	Set the metric value to give the redistributed routes (for any protocol except IGRP or EIGRP). The <i>metric value</i> is an integer from -294967295 to 294967295.
Step 18	set metric <i>bandwidth delay reliability loading mtu</i>	Set the metric value to give the redistributed routes (for IGRP or EIGRP only): <i>bandwidth</i>—Metric value or IGRP bandwidth of the route in kilobits per second in the range 0 to 4294967295 <i>delay</i>—Route delay in tens of microseconds in the range 0 to 4294967295. <i>reliability</i>—Likelihood of successful packet transmission expressed as a number between 0 and 255, where 255 means 100 percent reliability and 0 means no reliability. <i>loading</i>—Effective bandwidth of the route expressed as a number from 0 to 255 (255 is 100 percent loading). <i>mtu</i>—Minimum maximum transmission unit (MTU) size of the route in bytes in the range 0 to 4294967295.
Step 19	set metric-type { type-1 type-2 }	Set the OSPF external metric type for redistributed routes.
Step 20	set metric-type internal	Set the multi-exit discriminator (MED) value on prefixes advertised to external BGP neighbor to match the IGP metric of the next hop.
Step 21	set weight	Set the BGP weight for the routing table. The value can be from 1 to 65535.
Step 22	end	Return to privileged EXEC mode.
Step 23	show route-map	Display all route maps configured or only the one specified to verify configuration.
Step 24	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete an entry, use the **no route-map** *map tag* global configuration command or the **no match** or **no set** route-map configuration commands.

You can distribute routes from one routing domain into another and control route distribution.

Beginning in privileged EXEC mode, follow these steps to control route redistribution. Note that the keywords are the same as defined in the previous procedure.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router { bgp eigrp igrp isis iso-igrp ospf rip }	Enter router configuration mode.
Step 3	redistribute <i>protocol</i> [<i>process-id</i>] { level-1 level-1-2 level-2 } [metric <i>metric-value</i>] [metric-type <i>type-value</i>] [match internal external <i>type-value</i>] [tag <i>tag-value</i>] [route-map <i>map-tag</i>] [weight <i>weight</i>] [subnets]	Redistribute routes from one routing protocol to another routing protocol. If no route-maps are specified, all routes are redistributed. If the keyword route-map is specified with no <i>map-tag</i> , no routes are distributed.
Step 4	default-metric <i>number</i>	Cause the current routing protocol to use the same metric value for all redistributed routes (BGP, RIP and OSPF).
Step 5	default-metric <i>bandwidth delay reliability loading mtu</i>	Cause the IGRP or EIGRP routing protocol to use the same metric value for all non-IGRP redistributed routes.
Step 6	no default-information { in out }	Disable the redistribution of default information between IGRP processes, which is enabled by default.
Step 7	end	Return to privileged EXEC mode.
Step 8	show route-map	Display all route maps configured or only the one specified to verify configuration.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable redistribution, use the **no** form of the commands.

The metrics of one routing protocol do not necessarily translate into the metrics of another. For example, the RIP metric is a hop count, and the IGRP metric is a combination of five qualities. In these situations, an artificial metric is assigned to the redistributed route. Uncontrolled exchanging of routing information between different routing protocols can create routing loops and seriously degrade network operation.

If you have not defined a default redistribution metric that replaces metric conversion, some automatic metric translations occur between routing protocols:

- RIP can automatically redistribute static routes. It assigns static routes a metric of 1 (directly connected).
- IGRP can automatically redistribute static routes and information from other IGRP-routed autonomous systems. IGRP assigns static routes a metric that identifies them as directly connected. It does not change the metrics of routes derived from IGRP updates from other autonomous systems.
- Any protocol can redistribute other routing protocols if a default mode is in effect.

Configuring Policy-Based Routing

You can use policy-based routing (PBR) to configure a defined policy for traffic flows. By using PBR, you can have more control over routing by reducing the reliance on routes derived from routing protocols. PBR can determine and implement routing policies that allow or deny paths based on:

- Identity of a particular end system
- Application
- Protocol

You can use PBR to provide equal-access and source-sensitive routing, routing based on interactive versus batch traffic, or routing based on dedicated links. For example, you could transfer stock records to a corporate office on a high-bandwidth, high-cost link for a short time while transmitting routine application data such as e-mail over a low-bandwidth, low-cost link.

With PBR, you classify traffic using access control lists (ACLs) and then make traffic go through a different path. PBR is applied to incoming packets. All packets received on an interface with PBR enabled are passed through route maps. Based on the criteria defined in the route maps, packets are forwarded (routed) to the appropriate next hop.

- If packets do not match any route map statements, all set clauses are applied.
- If a statement is marked as deny, packets meeting the match criteria are sent through normal forwarding channels, and destination-based routing is performed.
- If a statement is marked as permit and the packets do not match any route-map statements, the packets are sent through the normal forwarding channels, and destination-based routing is performed.

For more information about configuring route maps, see the [“Using Route Maps to Redistribute Routing Information” section on page 28-90](#).

You can use standard IP ACLs to specify match criteria for a source address or extended IP ACLs to specify match criteria based on an application, a protocol type, or an end station. The process proceeds through the route map until a match is found. If no match is found, or if the route map is a deny, normal destination-based routing occurs. There is an implicit deny at the end of the list of match statements.

If match clauses are satisfied, you can use a set clause to specify the IP addresses identifying the next hop router in the path.

**Note**

For details about PBR commands and keywords, refer to the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

PBR Configuration Guidelines

Before configuring PBR, you should be aware of this information:

- To use PBR, you must first enable the PBR template by using the **sdm prefer routing-pbr** global configuration command. PBR is not supported with the routing, VLAN or default template. For more information on the SDM templates, see [Chapter 6, “Configuring SDM Templates.”](#)
- VRFs are not compatible with the PBR template. Therefore, multiprotocol label switching virtual private networks (MPLS VPNs) and multi-VRF CE are not compatible with PBR.
- Multicast traffic is not policy-routed. PBR applies to only to unicast traffic.
- You can enable PBR on a routed port or an SVI.

- You can apply a policy route map to an EtherChannel port channel in Layer 3 mode, but you cannot apply a policy route map to a physical interface that is a member of the EtherChannel. If you try to do so, the command is rejected. When a policy route map is applied to a physical interface, that interface cannot become a member of an EtherChannel.
- You can define a maximum of 246 IP policy route-maps on the switch.
- You can define a maximum of 512 access control entries (ACEs) for PBR on the switch.
- The number of TCAM entries used by PBR depends on the route map itself, the ACLs used, and the order of the ACLs and route-map entries.
- Policy-based routing based on packet length, IP precedence and TOS, set interface, set default next hop, or set default interface are not supported. Policy maps with no valid set actions or with set action set to *Don't Fragment* are not supported.

Enabling PBR

By default, PBR is disabled on the switch. To enable PBR, you must create a route map that specifies the match criteria and the resulting action if all of the match clauses are met. Then, you must enable PBR for that route map on an interface. All packets arriving on the specified interface matching the match clauses are subject to PBR.



Note

Before enabling PBR, configure the PBR template by entering the **sdm prefer routing-pbr** global configuration command and reloading the switch.

PBR can be fast-switched or implemented at speeds that do not slow down the switch. Fast-switched PBR supports most match and set commands. PBR must be enabled before you enable fast-switched PBR. Fast-switched PBR is disabled by default.

Packets that are generated by the switch, or local packets, are not normally policy-routed. When you globally enable local PBR on the switch, all packets that originate on the switch are subject to local PBR. Local PBR is disabled by default.

Beginning in privileged EXEC mode, follow these steps to configure PBR:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	route-map <i>map-tag</i> [permit deny] [<i>sequence number</i>]	Define any route maps used to control where packets are output, and enter route-map configuration mode. <i>map-tag</i> —A meaningful name for the route map. The ip policy route-map interface configuration command uses this name to reference the route map. Multiple route maps might share the same map tag name. (Optional) If permit is specified and the match criteria are met for this route map, the route is policy-routed as controlled by the set actions. If deny is specified, the route is not policy-routed. <i>sequence number</i> (Optional)— Number that shows the position of a new route map in the list of route maps already configured with the same name.

	Command	Purpose
Step 3	match ip address { <i>access-list-number</i> <i>access-list-name</i> } [... <i>access-list-number</i> ... <i>access-list-name</i>]	Match the source and destination IP address that is permitted by one or more standard or extended access lists. If you do not specify a match command, the route map applies to all packets.
Step 4	set ip next-hop <i>ip-address</i> [... <i>ip-address</i>]	Specify the action to take on the packets that match the criteria. Set next hop to which to route the packet (the next hop must be adjacent).
Step 5	exit	Return to global configuration mode.
Step 6	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to configure.
Step 7	ip policy route-map <i>map-tag</i>	Enable PBR on a Layer 3 interface, and identify the route map to use. You can configure only one route map on an interface. However, you can have multiple route map entries with different sequence numbers. These entries are evaluated in sequence number order until the first match. If there is no match, packets are routed as usual.
Step 8	ip route-cache policy	(Optional) Enable fast-switching PBR. You must first enable PBR before enabling fast-switching PBR.
Step 9	exit	Return to global configuration mode.
Step 10	ip local policy route-map <i>map-tag</i>	(Optional) Enable local PBR to perform policy-based routing on packets originating at the switch. This applies to packets generated by the switch and not to incoming packets.
Step 11	end	Return to privileged EXEC mode.
Step 12	show route-map [<i>map-name</i>]	(Optional) Display all route maps configured or only the one specified to verify configuration.
Step 13	show ip policy	(Optional) Display policy route maps attached to interfaces.
Step 14	show ip local policy	(Optional) Display whether or not local policy routing is enabled and, if so, the route map being used.
Step 15	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no route-map** *map-tag* global configuration command or the **no match** or **no set** route-map configuration commands to delete an entry. Use the **no ip policy route-map** *map-tag* interface configuration command to disable PBR on an interface. Use the **no ip route-cache policy** interface configuration command to disable fast-switching PBR. Use the **no ip local policy route-map** *map-tag* global configuration command to disable policy-based routing on packets originating on the switch.

Filtering Routing Information

You can filter routing protocol information by performing the tasks described in this section.



Note

When routes are redistributed between OSPF processes, no OSPF metrics are preserved.

Setting Passive Interfaces

To prevent other routers on a local network from dynamically learning about routes, you can use the **passive-interface** router configuration command to keep routing update messages from being sent through a router interface. When you use this command in the OSPF protocol, the interface address you specify as passive appears as a stub network in the OSPF domain. OSPF routing information is neither sent nor received through the specified router interface.

In networks with many interfaces, to avoid having to manually set them as passive, you can set all interfaces to be passive by default by using the **passive-interface default** router configuration command and manually setting interfaces where adjacencies are desired.

Beginning in privileged EXEC mode, follow these steps to configure passive interfaces:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router {bgp eigrp igmp isis iso-igmp ospf rip}	Enter router configuration mode.
Step 3	passive-interface interface-id	Suppress sending routing updates through the specified Layer 3 interface.
Step 4	passive-interface default	(Optional) Set all interfaces as passive by default.
Step 5	no passive-interface interface type	(Optional) Activate only those interfaces that need to have adjacencies sent.
Step 6	network network-address	(Optional) Specify the list of networks for the routing process. The <i>network-address</i> is an IP address.
Step 7	end	Return to privileged EXEC mode.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use a network monitoring privileged EXEC command such as **show ip ospf interface** to verify the interfaces that you enabled as passive, or use the **show ip interface** privileged EXEC command to verify the interfaces that you enabled as active.

To re-enable the sending of routing updates, use the **no passive-interface interface-id** router configuration command. The **default** keyword sets all interfaces as passive by default. You can then configure individual interfaces where you want adjacencies by using the **no passive-interface** router configuration command. The **default** keyword is useful in Internet service provider and large enterprise networks where many of the distribution routers have more than 200 interfaces.

Controlling Advertising and Processing in Routing Updates

You can use the **distribute-list** router configuration command with access control lists to suppress routes from being advertised in routing updates and to prevent other routers from learning one or more routes. When used in OSPF, this feature applies to only external routes, and you cannot specify an interface name.

You can also use a **distribute-list** router configuration command to avoid processing certain routes listed in incoming updates. (This feature does not apply to OSPF.)

Beginning in privileged EXEC mode, follow these steps to control the advertising or processing of routing updates:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router {bgp eigrp igmp isis iso-igrp ospf rip}	Enter router configuration mode.
Step 3	distribute-list { <i>access-list-number</i> <i>access-list-name</i> } out [<i>interface-name</i> <i>routing process</i> <i>autonomous-system-number</i>]	Permit or deny routes from being advertised in routing updates, depending upon the action listed in the access list.
Step 4	distribute-list { <i>access-list-number</i> <i>access-list-name</i> } in [<i>type-number</i>]	Suppress processing in routes listed in updates.
Step 5	end	Return to privileged EXEC mode.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no distribute-list in** router configuration command to change or cancel a filter. To cancel suppression of network advertisements in updates, use the **no distribute-list out** router configuration command.

Filtering Sources of Routing Information

Because some routing information might be more accurate than others, you can use filtering to prioritize information coming from different sources. An *administrative distance* is a rating of the trustworthiness of a routing information source, such as a router or group of routers. In a large network, some routing protocols can be more reliable than others. By specifying administrative distance values, you enable the router to intelligently discriminate between sources of routing information. The router always picks the route whose routing protocol has the lowest administrative distance. [Table 28-18 on page 28-89](#) shows the default administrative distances for various routing information sources.

Because each network has its own requirements, there are no general guidelines for assigning administrative distances.

Beginning in privileged EXEC mode, follow these steps to filter sources of routing information:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router {bgp eigrp igmp isis iso-igrp ospf rip}	Enter router configuration mode.
Step 3	distance weight { <i>ip-address</i> { <i>ip-address mask</i> }} [<i>ip access list</i>]	Define an administrative distance. <i>weight</i> —The administrative distance as an integer from 10 to 255. Used alone, <i>weight</i> specifies a default administrative distance that is used when no other specification exists for a routing information source. Routes with a distance of 255 are not installed in the routing table. (Optional) <i>ip access list</i> —An IP standard or extended access list to be applied to incoming routing updates.
Step 4	end	Return to privileged EXEC mode.

	Command	Purpose
Step 5	show ip protocols	Display the default administrative distance for a specified routing process.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a distance definition, use the **no distance** router configuration command.

Managing Authentication Keys

Key management is a method of controlling authentication keys used by routing protocols. Not all protocols can use key management. Authentication keys are available for EIGRP and RIP Version 2.

Before you manage authentication keys, you must enable authentication. See the appropriate protocol section to see how to enable authentication for that protocol. To manage authentication keys, define a key chain, identify the keys that belong to the key chain, and specify how long each key is valid. Each key has its own key identifier (specified with the **key number** key chain configuration command), which is stored locally. The combination of the key identifier and the interface associated with the message uniquely identifies the authentication algorithm and Message Digest 5 (MD5) authentication key in use.

You can configure multiple keys with life times. Only one authentication packet is sent, regardless of how many valid keys exist. The software examines the key numbers in order from lowest to highest, and uses the first valid key it encounters. The lifetimes allow for overlap during key changes. Note that the router must know these lifetimes.

Beginning in privileged EXEC mode, follow these steps to manage authentication keys:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	key chain <i>name-of-chain</i>	Identify a key chain, and enter key chain configuration mode.
Step 3	key number	Identify the key number. The range is 0 to 2147483647.
Step 4	key-string <i>text</i>	Identify the key string. The string can contain from 1 to 80 uppercase and lowercase alphanumeric characters, but the first character cannot be a number.
Step 5	accept-lifetime <i>start-time</i> { infinite <i>end-time</i> duration <i>seconds</i> }	(Optional) Specify the time period during which the key can be received. The <i>start-time</i> and <i>end-time</i> syntax can be either <i>hh:mm:ss Month date year</i> or <i>hh:mm:ss date Month year</i> . The default is forever with the default <i>start-time</i> and the earliest acceptable date as January 1, 1993. The default <i>end-time</i> and duration is infinite .
Step 6	send-lifetime <i>start-time</i> { infinite <i>end-time</i> duration <i>seconds</i> }	(Optional) Specify the time period during which the key can be sent. The <i>start-time</i> and <i>end-time</i> syntax can be either <i>hh:mm:ss Month date year</i> or <i>hh:mm:ss date Month year</i> . The default is forever with the default <i>start-time</i> and the earliest acceptable date as January 1, 1993. The default <i>end-time</i> and duration is infinite.

	Command	Purpose
Step 7	end	Return to privileged EXEC mode.
Step 8	show key chain	Display authentication key information.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the key chain, use the **no key chain *name-of-chain*** global configuration command.

Monitoring and Maintaining the IP Network

You can remove all contents of a particular cache, table, or database. You can also display specific statistics. Use the privileged EXEC commands in [Table 28-19](#) to clear routes or display status:

Table 28-19 Commands to Clear IP Routes or Display Route Status

Command	Purpose
clear ip route { <i>network</i> [<i>mask</i> *]}	Clear one or more routes from the IP routing table.
show ip protocols	Display the parameters and state of the active routing protocol process.
show ip route [<i>address</i> [<i>mask</i>] [longer-prefixes]] [<i>protocol</i> [<i>process-id</i>]]	Display the current state of the routing table.
show ip route summary	Display the current state of the routing table in summary form.
show ip route supernets-only	Display supernets.
show ip cache	Display the routing table used to switch IP traffic.
show route-map [<i>map-name</i>]	Display all route maps configured or only the one specified.



Configuring HSRP

This chapter describes how to use Hot Standby Router Protocol (HSRP) on the Catalyst 3750 Metro switch to provide routing redundancy for routing IP traffic without being dependent on the availability of any single router.



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 IOS IP and IP Routing Command Reference for Release 12.1*.

This chapter consists of these sections:

- [Understanding HSRP, page 29-1](#)
- [Configuring HSRP, page 29-3](#)
- [Displaying HSRP Configurations, page 29-10](#)

Understanding HSRP

HSRP is Cisco's standard method of providing high network availability by providing first-hop redundancy for IP hosts on an IEEE 802 LAN configured with a default gateway IP address. HSRP routes IP traffic without relying on the availability of any single router. It enables a set of router interfaces to work together to present the appearance of a single virtual router or default gateway to the hosts on a LAN. When HSRP is configured on a network or segment, it provides a virtual Media Access Control (MAC) address and an IP address that is shared among a group of configured routers. HSRP allows two or more HSRP-configured routers to use the MAC address and IP network address of a virtual router. The virtual router does not exist; it represents the common target for routers that are configured to provide backup to each other. One of the routers is selected to be the active router and another to be the standby router, which assumes control of the group MAC address and IP address should the designated active router fail.



Note

Routers in an HSRP group can be any router interface that supports HSRP, including Catalyst 3750 Metro switch routed ports and switch virtual interfaces (SVIs).


HSRP provides high network availability by providing redundancy for IP traffic from hosts on networks. In a group of router interfaces, the active router is the router of choice for routing packets; the standby router is the router that takes over the routing duties when an active router fails or when preset conditions are met.

HSRP is useful for hosts that do not support a router discovery protocol and cannot switch to a new router when their selected router reloads or loses power. When HSRP is configured on a network segment, it provides a virtual MAC address and an IP address that is shared among router interfaces in a group of router interfaces running HSRP. The router selected by the protocol to be the active router receives and routes packets destined for the group's MAC address. For n routers running HSRP, there are $n + 1$ IP and MAC addresses assigned.

HSRP detects when the designated active router fails, and a selected standby router assumes control of the Hot Standby group's MAC and IP addresses. A new standby router is also selected at that time. Devices running HSRP send and receive multicast UDP-based hello packets to detect router failure and to designate active and standby routers. When HSRP is configured on an interface, Internet Control Message Protocol (ICMP) redirect messages are disabled by default for the interface.

You can configure multiple Hot Standby groups among switches that are operating in Layer 3 to make more use of the redundant routers. To do so, specify a group number for each Hot Standby command group you configure for an interface. For example, you might configure an interface on switch 1 as an active router and one on switch 2 as a standby router and also configure another interface on switch 2 as an active router with another interface on switch 1 as its standby router.

Figure 29-1 shows a segment of a network configured for HSRP. Each router is configured with the MAC address and IP network address of the virtual router. Instead of configuring hosts on the network with the IP address of Router A, you configure them with the IP address of the virtual router as their default router. When Host C sends packets to Host B, it sends them to the MAC address of the virtual router. If for any reason, Router A stops transferring packets, Router B responds to the virtual IP address and virtual MAC address and becomes the active router, assuming the active router duties. Host C continues to use the IP address of the virtual router to address packets destined for Host B, which Router B now receives and sends to Host B. Until Router A resumes operation, HSRP allows Router B to provide uninterrupted service to users on Host C's segment that need to communicate with users on Host B's segment and also continues to perform its normal function of handling packets between the Host A segment and Host B.

Figure 29-1 Typical HSRP Configuration

101361

Configuring HSRP

These sections include HSRP configuration information:

- [Default HSRP Configuration, page 29-4](#)
- [HSRP Configuration Guidelines, page 29-4](#)
- [Enabling HSRP, page 29-5](#)
- [Configuring HSRP Group Attributes, page 29-6](#)

Default HSRP Configuration

Table 29-1 shows the default HSRP configuration.

Table 29-1 Default HSRP Configuration

Feature	Default Setting
HSRP groups	None configured
Standby group number	0
Standby MAC address	System assigned as: 0000.0c07.acXX, where XX is the HSRP group number
Standby priority	100
Standby delay	0 (no delay)
Standby track interface priority	10
Standby hello time	3 seconds
Standby holdtime	10 seconds

HSRP Configuration Guidelines

Follow these guidelines when configuring HSRP:

- HSRP can be configured on a maximum of 32 VLAN or routing interfaces.
- In the following procedures, the specified interface must be one of these Layer 3 interfaces:
 - Routed port: a physical port configured as a Layer 3 port by entering the **no switchport** interface configuration command.
 - SVI: a VLAN interface created by using the **interface vlan** *vlan_id* global configuration command and by default a Layer 3 interface.
 - Etherchannel port channel in Layer 3 mode: a port-channel logical interface created by using the **interface port-channel** *port-channel-number* global configuration command and binding the Ethernet interface into the channel group. For more information, see the “[Configuring Layer 3 EtherChannels](#)” section on page 27-12.
- All Layer 3 interfaces must have IP addresses assigned to them. See the “[Configuring Layer 3 Interfaces](#)” section on page 9-18.

Enabling HSRP

The **standby ip** interface configuration command activates HSRP on the configured interface. If an IP address is specified, that address is used as the designated address for the Hot Standby group. If no IP address is specified, the address is learned through the standby function. You must configure at least one routing port on the cable with the designated address. Configuring an IP address always overrides another designated address currently in use.

When the **standby ip** command is enabled on an interface and proxy ARP is enabled, if the interface's Hot Standby state is active, proxy ARP requests are answered using the Hot Standby group MAC address. If the interface is in a different state, proxy ARP responses are suppressed.



Note

When multi-VRF CE is configured, you cannot assign the same HSRP standby address to two different VPNs.

Beginning in privileged EXEC mode, follow these steps to create or enable HSRP on a Layer 3 interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the Layer 3 interface on which you want to enable HSRP.
Step 3	standby [<i>group-number</i>] ip [<i>ip-address</i> [<i>secondary</i>]]	Create (or enable) the HSRP group using its number and virtual IP address. (Optional) <i>group-number</i>—The group number on the interface for which HSRP is being enabled. The range is 0 to 255; the default is 0. If there is only one HSRP group, you do not need to enter a group number. (Optional on all but one interface) <i>ip-address</i>—The virtual IP address of the hot standby router interface. You must enter the virtual IP address for at least one of the interfaces; it can be learned on the other interfaces. (Optional) secondary—The IP address is a secondary hot standby router interface. If neither router is designated as a secondary or standby router and no priorities are set, the primary IP addresses are compared and the higher IP address is the active router, with the next highest as the standby router.
Step 4	end	Return to privileged EXEC mode.
Step 5	show standby [<i>interface-id</i> [<i>group</i>]]	Verify the configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no standby** [*group-number*] **ip** [*ip-address*] interface configuration command to disable HSRP.

This example shows how to activate HSRP for group 1 on an interface. The IP address used by the hot standby group is learned by using HSRP.

**Note**

This procedure is the minimum number of steps required to enable HSRP.

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# no switchport
Switch(config-if)# standby 1 ip
Switch(config-if)# end
Switch# show standby
```

Configuring HSRP Group Attributes

Although HSRP can run with no other configuration required, you can configure attributes for the HSRP group, including authentication, priority, preemption and preemption delay, timers, or MAC address.

Configuring HSRP Priority

The **standby priority**, **standby preempt**, and **standby track** interface configuration commands are all used to set characteristics for determining active and standby routers and behavior regarding when a new active router takes over. When configuring priority, follow these guidelines:

- Assigning priority helps select the active and standby routers. If preemption is enabled, the router with the highest priority becomes the designated active router. If priorities are equal, the primary IP addresses are compared, and the higher IP address has priority.
- The highest number (1 to 255) represents the highest priority (most likely to become the active router).
- When setting the priority, preempt, or both, you must specify at least one keyword (**priority**, **preempt**, or both).
- The priority of the device can change dynamically if an interface is configured with the **standby track** command and another interface on the router goes down.
- The **standby track** interface configuration command ties the router hot standby priority to the availability of its interfaces and is useful for tracking interfaces that are not configured for HSRP. When a tracked interface fails, the hot standby priority on the device on which tracking has been configured decreases by 10. If an interface is not tracked, its state changes do not affect the hot standby priority of the configured device. For each interface configured for hot standby, you can configure a separate list of interfaces to be tracked.
- The **standby track interface-priority** interface configuration command specifies how much to decrement the hot standby priority when a tracked interface goes down. When the interface comes back up, the priority is incremented by the same amount.
- When multiple tracked interfaces are down and *interface-priority* values have been configured, the configured priority decrements are cumulative. If tracked interfaces that were not configured with priority values fail, the default decrement is 10, and it is noncumulative.
- When routing is first enabled for the interface, it does not have a complete routing table. If it is configured to preempt, it becomes the active router, even though it is unable to provide adequate routing services. To solve this problem, configure a delay time to allow the router to update its routing table.

Beginning in privileged EXEC mode, use one or more of these steps to configure HSRP priority characteristics on an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the HSRP interface on which you want to set priority.
Step 3	standby [<i>group-number</i>] priority <i>priority</i> [preempt [delay <i>delay</i>]]	Set a priority value used in choosing the active router. The range is 1 to 255; the default priority is 100. The highest number represents the highest priority. • (Optional) <i>group-number</i>—The group number to which the command applies. • (Optional) preempt—Select so that when the local router has a higher priority than the active router, it assumes control as the active router. • (Optional) delay—Set to cause the local router to postpone taking over the active role for the shown number of seconds. The range is 0 to 36000 (1 hour); the default is 0 (no delay before taking over). Use the no form of the command to restore the default values.
Step 4	standby [<i>group-number</i>] [priority <i>priority</i>] preempt [delay <i>delay</i>]	Configure the router to preempt , which means that when the local router has a higher priority than the active router, it assumes control as the active router. • (Optional) <i>group-number</i>—The group number to which the command applies. • (Optional) priority—Enter to set or change the group priority. The range is 1 to 255; the default is 100. • (Optional) delay—Set to cause the local router to postpone taking over the active role for the number of seconds shown. The range is 0 to 36000 (1 hour); the default is 0 (no delay before taking over). Use the no form of the command to restore the default values.
Step 5	standby [<i>group-number</i>] track <i>type number</i> [<i>interface-priority</i>]	Configure an interface to track other interfaces so that if one of the other interfaces goes down, the device's Hot Standby priority is lowered. • (Optional) <i>group-number</i>—The group number to which the command applies. • <i>type</i>—Enter the interface type (combined with interface number) that is tracked. • <i>number</i>—Enter the interface number (combined with interface type) that is tracked. • (Optional) <i>interface-priority</i>—Enter the amount by which the hot standby priority for the router is decremented or incremented when the interface goes down or comes back up. The default value is 10.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify the configuration of the standby groups.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no standby** *[group-number]* **priority** *priority* [**preempt** *[delay delay]*] and **no standby** *[group-number]* [**priority** *priority*] **preempt** *[delay delay]* interface configuration commands to restore default priority, preempt, and delay values.

Use the **no standby** *[group-number]* **track** *type number* *[interface-priority]* interface configuration command to remove the tracking.

This example activates a port, sets an IP address and a priority of 120 (higher than the default value), and waits for 300 seconds (5 minutes) before attempting to become the active router:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# no switchport
Switch(config-if)# standby ip 172.19.108.254
Switch(config-if)# standby priority 120 preempt delay 300
Switch(config-if)# end
Switch#
```

Configuring HSRP Authentication and Timers

You can optionally configure an HSRP authentication string or change the hello-time interval and holdtime.

When configuring these attributes, follow these guidelines:

- The authentication string is sent unencrypted in all HSRP messages. You must configure the same authentication string on all routers and access servers on a cable to ensure interoperation. Authentication mismatch prevents a device from learning the designated Hot Standby IP address and timer values from other routers configured with HSRP.
- Routers or access servers on which standby timer values are not configured can learn timer values from the active or standby router. The timers configured on an active router always override any other timer settings.
- All routers in a Hot Standby group should use the same timer values. Normally, the *holdtime* is greater than or equal to 3 times the *hellotime*.

Beginning in privileged EXEC mode, use one or more of these steps to configure HSRP authentication and timers on an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the HSRP interface on which you want to set authentication.
Step 3	standby <i>[group-number]</i> authentication <i>string</i>	(Optional) authentication <i>string</i> —Enter a string to be carried in all HSRP messages. The authentication string can be up to eight characters in length; the default string is cisco . (Optional) <i>group-number</i> —The group number to which the command applies.

	Command	Purpose
Step 4	standby [<i>group-number</i>] timers <i>hellotime</i> <i>holdtime</i>	(Optional) Configure the time between hello packets and the time before other routers declare the active router to be down. <i>group-number</i>—The group number to which the command applies. <i>hellotime</i>—The hello interval in seconds. The range is from 1 to 255; the default is 3 seconds. <i>holdtime</i>—The time in seconds before the active or standby router is declared to be down. The range is from 1 to 255; the default is 10 seconds.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify the configuration of the standby groups.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no standby** [*group-number*] **authentication** *string* interface configuration command to delete an authentication string. Use the **no standby** [*group-number*] **timers** *hellotime* *holdtime* interface configuration command to restore timers to their default values.

This example shows how to configure *word* as the authentication string required to allow Hot Standby routers in group 1 to interoperate:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# no switchport
Switch(config-if)# standby 1 authentication word
Switch(config-if)# end
Switch#
```

This example shows how to set the timers on standby group 1 with the time between hello packets at 5 seconds and the time after which a router is considered down to be 15 seconds:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# no switchport
Switch(config-if)# standby 1 ip
Switch(config-if)# standby 1 timers 5 15
Switch(config-if)# end
Switch#
```

Displaying HSRP Configurations

From privileged EXEC mode, use this command to display HSRP settings:

show standby [*interface-id* [*group*]] [**brief**] [**detail**]

You can display HSRP information for the whole switch, for a specific interface, for an HSRP group, or for an HSRP group on an interface. You can also specify whether to display a concise overview of HSRP information or detailed HSRP information. The default display is **detail**. If there are a large number of HSRP groups, using the **show standby** command without qualifiers can result in an unwieldy display.

This is an example of output from the **show standby** privileged EXEC command, displaying HSRP information for two standby groups (group 1 and group 100):

```
Switch# show standby
VLAN1 - Group 1
  Local state is Standby, priority 105, may preempt
  Hellotime 3 holdtime 10
  Next hello sent in 00:00:02.182
  Hot standby IP address is 10.0.0.1 configured
  Active router is 172.20.138.35 expires in 00:00:09
  Standby router is local
  Standby virtual mac address is 0000.0c07.ac01
  Name is bbb
VLAN1 - Group 100
  Local state is Active, priority 105, may preempt
  Hellotime 3 holdtime 10
  Next hello sent in 00:00:02.262
  Hot standby IP address is 172.20.138.51 configured
  Active router is local
  Standby router is unknown expired
  Standby virtual mac address is 0000.0c07.ac64
  Name is test
```



Configuring MPLS and EoMPLS

This chapter describes how to configure multiprotocol label switching (MPLS) and Ethernet over MPLS (EoMPLS) on the Catalyst 3750 Metro switch. MPLS is a packet-switching technology that integrates link layer (Layer 2) switching with network layer (Layer 3) routing. With MPLS, data is transferred over any combination of Layer 2 technologies, using any Layer 3 protocol, with increased scalability. MPLS supports different routes between a source and destination over a router-based Internet backbone.

EoMPLS is tunneling mechanism that transports Layer 2 Ethernet frames over an MPLS network. You can connect two Layer 2 networks that are in different locations, without requiring bridges, routers, or switches at the locations. You enable the MPLS backbone to accept Layer 2 traffic by configuring the label-edge routers (LERs) at both ends of the MPLS backbone.

MPLS functionality is supported only on the enhanced-services (ES) ports; EoMPLS is supported on standard and ES ports.



Note

For more information about MPLS, refer to the *Cisco IOS Switching Services Configuration Guide for Release 12.1*. For complete syntax and usage information for the MPLS commands used in this chapter, refer to the *Cisco IOS Switching Services Command Reference for Release 12.1*. For more information about EoMPLS commands, refer also to the command reference for this release.

This chapter contains these sections:

- [Understanding MPLS Services, page 30-1](#)
- [Understanding MPLS VPNs, page 30-2](#)
- [Configuring MPLS VPNs, page 30-6](#)
- [Understanding EoMPLS, page 30-12](#)
- [Enabling EoMPLS, page 30-15](#)
- [Configuring MPLS and EoMPLS QoS, page 30-18](#)
- [Monitoring and Maintaining MPLS and EoMPLS, page 30-22](#)

Understanding MPLS Services

In conventional Layer 3 forwarding, as a packet travels across the network, each router extracts the packet-forwarding information from the Layer 3 header and uses this information as an index for a routing table lookup to determine the packet's next hop. In most cases, the only relevant field in the header is the destination address field, but in some cases other header fields are also relevant. For this reason, each router through which the packet passes must analyze the packet header.

With MPLS, the Layer 3 header is analyzed only once and then is mapped into a fixed-length, unstructured value called a *label*. Many different headers can map to the same label, as long as those headers always result in the same choice of next hop. In effect, a label represents a *forwarding-equivalence class*—that is, a set of packets that can be very different but that are indistinguishable to the forwarding function.

The initial choice of label can be based exclusively on the contents of the Layer 3 header, or it can be based on policy, allowing forwarding decisions at subsequent hops to be based on policy. After a label is chosen, a short label header is put at the front of the Layer 3 packet and carried across the network as part of the packet. At subsequent hops through each MPLS router in the network, labels are exchanged, and the router uses MPLS forwarding-table lookups for the label to make forwarding decisions. It is not necessary to re-analyze the packet header. Because the label is a fixed length and unstructured, the MPLS forwarding-table lookup process is straightforward and fast.

Each label-switching router (LSR) in the network makes an independent, local decision as to which label value is used to represent which forwarding equivalence class. This association is known as a *label binding*. Each LSR informs its neighbors of the label bindings that it has made. When a labeled packet is sent from LSR A to neighboring LSR B, the label value carried by the packet is the label value that B assigned to represent the packet's forwarding equivalence class. Thus, the label value changes as the IP packet travels through the network.

**Note**

Because the Catalyst 3750 Metro switch is used as a service-provider edge (PE) device, rather than a service-provider core router, it does not normally operate as an LSR. The switch only performs label switching when it is connected to two different provider core routers over the ES ports to provide a redundant path. In this case, the switch uses QoS policies to classify MPLS packets on egress for label switching.

A label represents a forwarding-equivalence class, but it does not represent a particular path through the network. In general, the path through the network continues to be chosen by the existing Layer 3 routing protocols, such as Open Shortest Path First (OSPF), Enhanced Interior Gateway Protocol (EIGRP), Intermediate-System-to-Intermediate-System (IS-IS), and Border Gateway Protocol (BGP). At each hop when a label is looked up, the choice of the next hop is determined by the dynamic routing algorithm.

Understanding MPLS VPNs

Using MPLS virtual private networks (VPNs) provides the capability to deploy and administer scalable Layer 3 VPN backbone services to business customers. A VPN is a secure IP-based network that shares resources on one or more physical networks. A VPN contains geographically dispersed sites that can communicate securely over a shared backbone.

VPN routes are distributed over the MPLS network by using multiprotocol BGP (MP-BGP), which also distributes the labels associated with each VPN route. MPLS VPN depends on VPN routing and forwarding (VRF) support to isolate the routing domains from each other. When routes are learned over an MPLS VPN, the switch learns the new route as a normal VRF route, except that the destination MAC address for the next hop is not the real address, but a specially formed address that contains an identifier that is allocated for the route. When an MPLS-VPN packet is received on a port, the switch looks up the labels in the routing table to determine what to do with the packet.

Each VPN is associated with one or more VPN VRF instances. A VRF includes routing and forwarding tables and rules that define the VPN membership of customer devices attached to the customer edge (CE) device. A customer site can be a member of multiple VPNs; however, a site can associate with only one VRF. A VRF has these elements:

- An IP routing table
- A Cisco Express Forwarding (CEF) table
- A set of interfaces that use the CEF forwarding table
- A set of rules and routing protocol parameters to control the information in the routing tables

A customer-site VRF contains all the routes available to the site from the VPNs to which it belongs. VPN routing information is stored in the IP routing table and the CEF table for each VRF. A separate set of tables is maintained for each VRF, which prevents information from being forwarded outside a VPN and prevents packets that are outside a VPN from being forwarded to a router within the VPN. Based on the routing information stored in the VRF IP routing table and the VRF CEF table, packets are forwarded to their destinations.

A PE router binds a label to each customer prefix that is learned from a CE device and includes the label in the network reachability information for the prefix that it advertises to other PE routers. When a PE router forwards a packet that is received from a CE device across the provider network, it labels the packet with the label learned from the destination PE router. When the destination PE router receives the labeled packet, it examines the label and uses it to direct the packet to the correct CE device. A customer data-packet carries two levels of labels when traversing the backbone:

- The top label directs the packet to the correct PE router.
- The second label defines how that PE router should forward the packet to the CE device.

VPN Benefits


MPLS VPNs allow service providers to deploy scalable VPNs and build the foundation to deliver value-added services, including:

- Connectionless service—MPLS VPNs are connectionless, which means that no prior action is required to establish communication between hosts. A connectionless VPN does not require tunnels and encryption for network privacy.
- Centralized service—MPLS VPNs are seen as private intranets, which allows delivery of targeted IP services to a group of users represented by a VPN.
- Scalability—MPLS-based VPNs use the peer model and Layer 3 connectionless architecture to leverage a highly scalable solution. The peer model requires a customer site to act as a peer to one PE router as opposed to all other customer provider-edge or CE devices that are members of the VPN. The PE routers maintain VPN routes for those VPNs who are members. Routers in the core network do not maintain any VPN routes.
- Security—MPLS VPNs offer the same level of security as connection-oriented VPNs. Packets from one VPN do not inadvertently go to another VPN. Security provided at the edge of a provider network ensures that packets received from a customer are placed on the correct VPN; security provided at the backbone ensures that VPN traffic is kept separate.
- Easy to create—Because MPLS VPNs are connectionless, no specific point-to-point connection maps or topologies are required, and you can add sites to intranets and extranets to form closed user groups.

- Flexible addressing—Customers can continue to use their present address spaces without network address translation (NAT) because the MPLS VPN provides a public and private view of the address. A NAT is required only if two VPNs with overlapping address spaces want to communicate.
- Straightforward migration—You can build MPLS VPNs over multiple network architectures. Migration for the end customer is simplified because the CE router is not required to support MPLS, so no customer's intranet modifications are needed.
- MPLS VPN also provides increased BGP functionality.

Figure 30-1 shows an example of a VPN with a service-provider backbone network, PE routers, and CE devices.

Figure 30-1 VPNs with a Service-Provider Backbone




Each VPN contains customer devices attached to the CE devices. The customer devices use VPNs to exchange information between devices, and the provider routers (P) are not aware of the VPNs.

Figure 30-2 shows five customer sites communicating within three VPNs. The VPNs can communicate with these sites:

VPN1: Sites 2 and 4

VPN2: Sites 1, 3, and 4

VPN3: Sites 1, 3, and 5

Figure 30-2 Customer Sites with VPNs

Distribution of VPN Routing Information

The distribution of VPN routing information is controlled through the use of VPN route target communities, implemented by BGP extended communities. VPN routing information is distributed in this manner:

- When a VPN route learned from a CE device is added to the BGP process, a list of VPN route target extended community attributes is associated with it. The attribute values are obtained from an export list of route targets associated with the VRF from which the route was learned.
- An import list of route target extended communities is also associated with each VRF. The import list defines route target extended community attributes that a route must have in order for the route to be imported into the VRF. For example, if the import list for a particular VRF includes route target communities A, B, and C, then any VPN route that carries any of those route target extended communities—A, B, or C—is imported into the VRF.

A PE router can learn an IP prefix from a CE device by static configuration, through a BGP session with the CE device, or through the routing information protocol (RIP) exchange with the CE router. The IP prefix is a member of the IPv4 address family. After it learns the IP prefix, the PE converts it into a VPN-IPv4 prefix by combining it with an 8-byte route distinguisher (RD). The generated prefix is a member of the VPN-IPv4 address family and uniquely identifies the customer address, even if the customer site is using globally nonunique (unregistered private) IP addresses.

BGP distributes reachability information for VPN-IPv4 prefixes for each VPN. BGP communication takes place at two levels: within IP domains, known as autonomous systems (internal BGP or IBGP), and between autonomous systems (external BGP or EBGp). PE-to-PE sessions are IBGP sessions, and PE-CE sessions are EBGp sessions.

BGP propagates reachability information for VPN-IPv4 prefixes among PE routers by using the BGP multiprotocol extensions, which define support for address families other than IPv4. It does this in a way that ensures that the routes for a given VPN are learned only by other members of that VPN, which enables members of the VPN to communicate with each other.

Configuring MPLS VPNs

This section includes this information about configuring MPLS VPNs on a Catalyst 3750 Metro switch used as a PE router:

- [Default MPLS Configuration, page 30-6](#)
- [MPLS VPN Configuration Guidelines, page 30-6](#)

These sections describe the required tasks:

- [Enabling MPLS, page 30-7](#)
- [Defining VPNs, page 30-8](#)
- [Configuring BGP Routing Sessions, page 30-9](#)
- [Configuring PE-to-PE Routing Sessions, page 30-9](#)

You must also configure a PE to CE routing session. These sections provide example configurations:

- [Configuring BGP PE-to-CE Routing Sessions, page 30-10](#)
- [Configuring RIP PE-to-CE Routing Sessions, page 30-10](#)
- [Configuring Static Route PE-to-CE Routing Sessions, page 30-11](#)

For an example of packet flow in an MPLS VPN, see the [“Packet Flow in an MPLS VPN” section on page 30-11](#).

Default MPLS Configuration

By default, label switching of IPv4 packets along normally routed paths is globally enabled. MPLS forwarding of IPv4 packets is disabled by default on interfaces.

If no distribution protocol is explicitly configured by the **mpls label protocol** global configuration command, tag distribution protocol (TDP) is the default label distribution protocol for the switch. Cisco recommends that you configure label distribution protocol (LDP) for MPLS.

If no protocol is explicitly configured for an interface, the default label distribution protocol for the switch is used. By default, the labels of all destinations are advertised to all LDP neighbors.

No VRFs are defined. The default routing table for an interface is the global routing table.

MPLS VPN Configuration Guidelines

MPLS requires that CEF is enabled on the switch. CEF is enabled by default. For more information about CEF, refer to the [“Configuring Cisco Express Forwarding” section on page 28-86](#).

The switch must connect to the MPLS network through an ES port. MPLS configuration is supported only on ES ports.

Do not configure VLAN mapping on an interface configured for MPLS.

The switch supports a total of 26 VRFs and VPNs.

VRFs are not compatible with the PBR template. If you configure the PBR template by entering the **sdm prefer routing-pbr** command, any preconfigured VRFs are removed from the configuration. PBR and VRFs cannot function on the same switch.

Enabling MPLS

To use MPLS in a network, such as the one shown in [Figure 30-1](#), MPLS must be globally enabled and explicitly configured on the PE routers.

Beginning in privileged EXEC mode, follow these steps to incrementally deploy MPLS through a network, assuming that packets to all destination prefixes should be label-switched:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip cef	Enable CEF on the device if it has been disabled.
Step 3	mpls ip	Enable MPLS forwarding of IPv4 packets on normally routed paths if it has been disabled.
Step 4	mpls label protocol ldp	Set the label protocol on the switch to LDP. The default protocol is TDP.
Step 5	mpls ldp advertise-labels [for <i>prefix-access-list</i> [to <i>peer-access-list</i>]]	Enable MPLS label advertising on the switch. If no keywords are included, there are no restrictions on which labels are advertised. (Optional) for <i>prefix-access-list</i>—Specify which destinations should have their labels advertised. (Optional) to <i>peer-access-list</i>—Specify which LDP neighbors should receive label advertisements. A label switch router (LSR) is identified by its router ID, which consists of the first 4 bytes of its 6-byte LDP identifier.
Step 6	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 ES interface connected to the MPLS network.
Step 7	mpls ip	Enable MPLS forwarding of IPv4 packets along normally routed paths for the interface.
Step 8	end	Return to privileged EXEC mode.
Step 9	show mpls forwarding table show mpls interfaces	Verify the configuration.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Repeat these steps for every PE router in the network and the appropriate interfaces until all routers and connected interfaces are enabled for MPLS.

Use the **no mpls ip** global configuration command to disable MPLS on the switch. Use the **no mpls label protocol ldp** global configuration command to return to the default TDP.

Defining VPNs

Beginning in privileged EXEC mode, follow these steps to define VPN routing instances on the PE router:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip routing	Enable IP routing (required only if routing is disabled).
Step 3	ip vrf <i>vrf-name</i>	Enter VRF configuration mode, and define the VPN routing instance by assigning a VRF name.
Step 4	rd <i>route-distinguisher</i>	Create a VRF table by specifying a route distinguisher. Enter either an AS number and an arbitrary number (xxx:y) or an IP address and arbitrary number (A.B.C.D:y).
Step 5	route-target { export import both } <i>route-target-ext-community</i>	Create a list of import, export, or import and export route target communities for the specified VRF. Enter either an AS system number and an arbitrary number (xxx:y) or an IP address and an arbitrary number (A.B.C.D:y). The <i>route-target-ext-community</i> should be the same as the <i>route-distinguisher</i> entered in Step 4.
Step 6	import map <i>route-map</i>	(Optional) Associate the specified import route map with the VRF.
Step 7	export map <i>route-map</i>	(Optional) Associate the specified export route map with the VRF.
Step 8	exit	Return to global configuration mode.
Step 9	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 ES or VLAN interface to be associated with the VRF.
Step 10	ip vrf forwarding <i>vrf-name</i>	Associate the Layer 3 interface with the VRF.
Step 11	end	Return to privileged EXEC mode.
Step 12	show ip vrf	Display the defined VRFs and interfaces.
Step 13	show ip route vrf show ip cef vrf <i>vrf-name</i>	Display the IP routing table for a VRF. Display the CEF forwarding table associated with a VRF.
Step 14	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip vrf** *vrf-name* global configuration command to delete a VRF and remove all interfaces from it. Use the **no ip vrf forwarding** interface configuration command to remove an interface from a VRF.

Configuring BGP Routing Sessions

Beginning in privileged EXEC mode, follow these steps on the PE router to configure BGP routing sessions in a provider network:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip routing	Enable IP routing (required only if routing is disabled).
Step 3	router bgp <i>autonomous-system-number</i>	Enable a BGP routing process, assign it the AS number passed to the other BGP routers, and enter router configuration mode. The AS number can be from 1 to 65535, with 64512 to 65535 designated as private autonomous numbers.
Step 4	neighbor { <i>ip-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i>	Specify a neighbor IP address or BGP peer group that identifies it to the local autonomous system. The AS number can be from 1 to 65535.
Step 5	neighbor <i>ip-address</i> activate	Activate the advertisement of the IPv4 address family.
Step 6	end	Return to privileged EXEC mode.
Step 7	show ip bgp neighbor	Verify the configuration.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no router bgp** *autonomous-system* global configuration command to delete the BGP routing session.

Configuring PE-to-PE Routing Sessions

Beginning in privileged EXEC mode, follow these steps on the PE router to configure a PE-to-PE routing session in a provider network that uses IBGP:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system-number</i>	Enter router configuration mode.
Step 3	address-family vpnv4 [unicast]	Enter address family configuration mode to configure routing sessions that use standard VPNv4 address prefixes. (Optional) unicast —Specify VPNv4 unicast address prefixes.
Step 4	neighbor <i>ip-address</i> remote-as <i>as-number</i>	Define an IBGP session between the PE routers.
Step 5	neighbor <i>ip-address</i> activate	Activate the advertisement of the IPv4 address family.
Step 6	end	Return to privileged EXEC mode.
Step 7	show ip bgp [ipv4] [neighbors] [vpnv4]	Verify BGP configuration. Display information about all BGP IPv4 prefixes.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no router bgp** *autonomous-system* global configuration command to delete the BGP routing session.

Configuring BGP PE-to-CE Routing Sessions

Beginning in privileged EXEC mode, follow these steps on the PE router to configure a BGP PE-to-CE routing session:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system-number</i>	Configure the BGP routing process with the AS number passed to other BGP routers, and enter router configuration mode.
Step 3	address-family ipv4 [unicast] vrf <i>vrf-name</i>	Define EGP parameters for PE-to-CE routing sessions, and enter VRF address-family configuration mode. Note The default is <i>off</i> for auto-summary and synchronization in the VRF address-family configuration mode.
Step 4	neighbor <i>address</i> remote-as <i>as-number</i>	Define an EBGP session between PE and CE routers.
Step 5	neighbor <i>address</i> activate	Activate the advertisement of the IPv4 address family.
Step 6	end	Return to privileged EXEC mode.
Step 7	show ip bgp [ipv4] [neighbors]	Verify BGP configuration. Display information about all BGP IPv4 prefixes.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no router bgp** *autonomous-system-number* global configuration command to delete the BGP routing session.

Configuring RIP PE-to-CE Routing Sessions

Beginning in privileged EXEC mode, follow these steps on the PE router to configure RIP PE-to-CE routing:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router rip	Enable RIP routing, and enter router configuration mode.
Step 3	address-family ipv4 [unicast] vrf <i>vrf-name</i>	Define RIP parameters for PE to CE routing sessions, and enter VRF address-family configuration mode. Note The default is <i>off</i> for auto-summary and synchronization in the VRF address-family configuration mode.
Step 4	network <i>prefix</i>	Enable RIP on the PE-to-CE link.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip rip database [<i>network-prefix</i>]	Verify the configuration.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no router rip** global configuration command to disable RIP routing.

Configuring Static Route PE-to-CE Routing Sessions


Beginning in privileged EXEC mode, follow these steps on the PE router to configure static routing:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip route vrf vrf-name prefix mask	Define the VRF static routing table to use for PE-to-CE sessions.
Step 3	router bgp autonomous-system-number	Enter the BGP routing process AS number, and enter router configuration mode.
Step 4	address-family ipv4 [unicast] vrf vrf-name	Define static route parameters for every PE-to-CE routing session, and enter VRF address-family configuration mode. Note The default is <i>off</i> for auto-summary and synchronization in the VRF address-family configuration mode.
Step 5	redistribute static	Redistribute VRF static routes into the VRF BGP table.
Step 6	redistribute connected	Redistribute directly connected networks into the VRF BGP table.
Step 7	end	Return to privileged EXEC mode.
Step 8	show ip bgp [ipv4]	Verify the configuration.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Packet Flow in an MPLS VPN

Figure 30-3 is an example of packet flow between two customer sites in an MPLS VPN network.

Figure 30-3 Sample MPLS VPN Packet Flow



101099

A customer (Fast Ethernet) port on switch PE1 is configured for routed operation in a VPN. The port uses static routing or a routing protocol (RIP, OSPF, EIGRP, or BGP) to forward packets. MP-BGP is configured over the PE1 switch ES port with a route distinguisher (RD) that is associated with the customer's VPN. MP-BGP is configured to redistribute the routes and their associated VPN labels over the ES port that is using this RD.

The packet flow follows these steps:

-
Step 1	Provider-edge switch PE1 (which could be a Catalyst 3750 Metro switch) receives a packet from the customer switch at site 1. The switch determines from the lookup table that the VRF is a VLAN running MPLS and uses the MPLS lookup table to determine what to do with the packet. The MPLS lookup table contains the peer LSR as the destination MAC address and the local interface as the source MAC address.
Step 2	PE1 finds a BGP route with the appropriate next hop and labels, adds the appropriate labels to the packet, and forwards the packet out of the ES port to the next hop router (P3).
Step 3	The P3 router receives the packet and forwards it over the MPLS-VPN network, based on the packet's top label—the interior gateway protocol (IGP) label—and then removes the top label.
Step 4	PE3 receives the packet, removes the MPLS encapsulation, and forwards the packet by using the VRF interface associated with the VPN label contained in the packet that has the customer-edge switch CE2 as the destination.
-

Understanding EoMPLS

Any Transport over MPLS (AToM) is a solution for transporting Layer 2 packets over an MPLS network, allowing service providers to use the MPLS network to provide connectivity between customer sites with existing Layer 2 networks. Instead of separate networks with network management environments, service providers can use the MPLS network to transport all types of traffic for different customers. The Catalyst 3750 Metro switch supports EoMPLS, a subset of AToM that uses a tunneling mechanism to carry Layer 2 Ethernet traffic. EoMPLS encapsulates Ethernet frames in MPLS packets and forwards them across the MPLS network. Each frame is transported as a single packet, and the PE routers connected to the backbone add and remove labels as appropriate for packet encapsulation:

- The ingress PE router receives an Ethernet frame and encapsulates the packet by removing the preamble, the start of frame delimiter (SFD), and the frame check sequence (FCS). The rest of the packet header is not changed.
- The ingress PE router adds a point-to-point virtual connection (VC) label and a label switched path (LSP) tunnel label for normal MPLS routing through the MPLS backbone.
- The network core routers use the LSP tunnel label to move the packet through the MPLS backbone and do not distinguish Ethernet traffic from any other types of packets in the MPLS backbone.
- At the other end of the MPLS backbone, the egress PE router receives the packet and de-encapsulates the packet by removing the LSP tunnel label if one is present. The PE router also removes the VC label from the packet.
- The PE router updates the header, if necessary, and sends the packet out the appropriate interface to the destination switch.

The MPLS backbone uses the tunnel labels to transport the packet between the PE routers. The egress PE router uses the VC label to select the outgoing interface for the Ethernet packet. EoMPLS tunnels are unidirectional; for bidirectional EoMPLS, you need to configure one tunnel in each direction.

The point-to-point VC requires you to configure VC endpoints at the two PE routers. Only the PE routers at the ingress and egress points of the MPLS backbone know about the VCs dedicated to transporting Layer 2 traffic. Other routers do not have table entries for these VCs.

This section includes additional information about these topics:

- [Interaction with Other Features, page 30-13](#)
- [EoMPLS Limitations, page 30-14](#)

Interaction with Other Features

This section describes how EoMPLS interacts other features. It includes these sections:


- [EoMPLS and 802.1Q Tunneling, page 30-13](#)
- [EoMPLS and Layer 2 Tunneling, page 30-14](#)
- [EoMPLS and QoS, page 30-14](#)

EoMPLS and 802.1Q Tunneling

IEEE 802.1Q tunneling enables service providers to use a single VLAN to support customers who have multiple VLANs, while preserving customer VLAN IDs and segregating traffic in different VLANs. For more information about 802.1Q tunneling, see [Chapter 13, “Configuring IEEE 802.1Q and Layer 2 Protocol Tunneling.”](#)

Figure 30-4 is an example configuration where 802.1Q-tunneled traffic is forwarded using EoMPLS over an MPLS network. To support 802.1Q tunneling in a topology where a Layer 2 device connects to an MPLS network through a switch functioning as a PE, the ingress LAN port on the PE that receives the 802.1Q tunnel-encapsulated traffic (PE1) is configured as a tunnel port that accepts VLAN 100 traffic. On PE1, the interface is configured for port-based EoMPLS forwarding, with PE2 as the destination IP address. When packets from VLANs 10 to 50 arrive from CE1, they are encapsulated in VLAN 100 and sent to the PE1 egress port that is connected to the MPLS network. At the egress port, an MPLS tag is added to the frame header before it is mapped to a VC and forwarded to the next MPLS PE (PE2).

Figure 30-4 EoMPLS Example



By entering the **mpls l2transport route** or the **xconnect** interface configuration command on either a VLAN for VLAN-based EoMPLS or an Ethernet port for port-based EoMPLS, you can configure an EoMPLS tunnel to forward traffic based on either the customer VLAN or the Ethernet port.

- To forward 802.1Q tunnel-encapsulated traffic through the MPLS core to a specific recipient on the other side of the MPLS network, configure port-based EoMPLS.
- To forward 802.1Q tunnel-encapsulated traffic from an access device to a PE router, configure VLAN-based EoMPLS.

EoMPLS and Layer 2 Tunneling

Layer 2 protocol tunneling over an EoMPLS link allows CDP, STP, and VTP protocol data units (PDUs) to be tunneled through an MPLS network. To support Layer 2 protocol tunneling when the Layer 2 device connects to an MPLS network through a switch functioning as a PE, you configure the ingress port on the PE that receives the Layer 2 protocol traffic as a tunnel port. The Layer 2 protocol traffic is encapsulated before it is forwarded over the MPLS network. For more information about Layer 2 protocol tunneling, see [Chapter 13, “Configuring IEEE 802.1Q and Layer 2 Protocol Tunneling.”](#)

EoMPLS and QoS

EoMPLS supports QoS by using three experimental bits in a label to determine the priority of packets. To support QoS between label edge routers (LERs), you set the experimental bits in both the VC and tunnel labels. EoMPLS QoS classification occurs on ingress, and you can only match on Layer 3 parameters (such as IP or DSCP), not Layer 2 parameters (CoS). See the [“Configuring MPLS and EoMPLS QoS” section on page 30-18](#) for more information about EoMPLS and QoS.

EoMPLS Limitations

These restrictions apply to EoMPLS:

- EoMPLS requires that at least one of the two ES ports be configured for MPLS. Therefore, if you want to run EoMPLS on an ES port, you can only configure it on the ES port that is not configured for MPLS.
- MTU—EoMPLS does not support packet fragmentation and reassembly. Therefore, the maximum transmission unit (MTU) of all intermediate links between endpoints must be sufficient to carry the largest Layer 2 VLAN received. The ingress and egress PE routers must have the same MTU value.
- Address Format—All loopback addresses on PE routers must be configured with 32-bit masks to ensure proper operation of MPLS forwarding. OSPF requires the use of loopback addresses.
- Packet Format—EoMPLS supports VLAN packets that conform to the IEEE 802.1Q standard. ISL encapsulation is not supported between PE and CE routers.
- The maximum number of VLANs using EoMPLS on a switch is 1005.
- Layer 2 connection restrictions:
 - You cannot have a direct Layer 2 connection between PEs with EoMPLS.
 - You cannot have more than one Layer 2 connection between routers if those routers are configured to transport Ethernet VLANs over the MPLS backbone. Adding a second Layer 2 connection causes the spanning-tree state to constantly toggle if you disable spanning tree on the peer router.

- EoMPLS and trunking restrictions:
 - To support Ethernet spanning-tree bridge protocol data units (BPDUs) across an EoMPLS backbone, you must disable spanning tree for the EoMPLS VLAN. This ensures that the EoMPLS VLANs are carried only on the trunk to the customer switch.
 - The native VLAN of a trunk cannot be an EoMPLS VLAN.
- You can enable EoMPLS on 802.1Q interfaces by using the **mpls l2transport route** or the **xconnect** interface configuration command.
- Do not configure VLAN mapping on an interface configured for EoMPLS.

Enabling EoMPLS

This section includes this information about configuring EoMPLS on a switch used as a PE router:

- [Default EoMPLS Configuration, page 30-15](#)
- [EoMPLS Configuration Guidelines, page 30-15](#)
- [Configuring EoMPLS, page 30-16](#)
- [Packet Flow in an EoMPLS Network, page 30-17](#)

Default EoMPLS Configuration

By default, EoMPLS is not configured.

The **mpls ldp router-id** command is disabled. No virtual connections are configured.

EoMPLS Configuration Guidelines

When you configure EoMPLS, you must follow these guidelines:

- EoMPLS requires that at least one of the two ES ports be configured for MPLS. Therefore, if you want to run EoMPLS on an ES port, you can only configure it on the ES port that is not configured for MPLS.
- Before enabling EoMPLS, you must enable dynamic MPLS labeling by using the **mpls ip** interface configuration command on all paths between the imposition and disposition LERs. MPLS is globally enabled by default.
- For VLAN-based EoMPLS, you must configure VLANs on the switch.
- EoMPLS operation between two PE routers requires an LDP session between the routers. The IP address used by each router as its LDP router ID must be IP-reachable by the other. Use the optional **mpls ldp router-id** global configuration command to control the selection of the LDP router ID by specifying the interface whose IP address should be used.
 - If the specified interface is up and has an IP address, you can use the command without the optional **force** keyword. When the router ID is selected, that IP address is selected as the router ID.
 - If the specified interface is not up or does not have an IP address, use the **force** keyword with the command to ensure that the IP address of the specified interface is used when that interface is brought up.

- Both PE routers require a loopback address that you can use to create a VC between the routers. When you use OSPF as the interior gateway protocol, you must configure all loopback addresses on PE routers with 32-bit masks to ensure proper operation of MPLS forwarding between the PE routers.
- Do not configure EoMPLS on an interface configured for VLAN mapping.

Configuring EoMPLS

You configure VLAN-based EoMPLS on a VLAN interface. When VLAN-based EoMPLS is enabled, the switch associates the tunnel and VC labels based on the VLAN ID. You use the same commands to enable port-based EoMPLS on an ES interface.

Beginning in privileged EXEC mode, follow these steps on the PE routers to configure EoMPLS to transport Layer 2 packets between two endpoints:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mpls label protocol ldp	Enable LDP for all interfaces. By default, TDP is enabled. This command causes all interfaces to use LDP.
Step 3	interface loopback0	Enter interface configuration mode for a loopback interface.
Step 4	ip address <i>ip-address subnet mask</i>	Assign an IP address to the loopback interface.
Step 5	exit	Return to global configuration mode
Step 6	mpls ldp router-id loopback0 force	(Optional) Force the IP address of loopback interface 0 to be used as the router ID.
Step 7	interface <i>interface-id</i>	Enter a Layer 3 VLAN (for VLAN-based EoMPLS) or the interface ID of an ES port (for port-based EoMPLS), and enter interface configuration mode.
Step 8	mpls ip	Enable MPLS on the interface.
Step 9	mpls l2transport route <i>destination vc-id</i> or xconnect <i>destination vc-id</i>	Configure the interface to transport the Layer 2 VLAN packets over MPLS. • <i>destination</i>—The IP address of the PE router at the other end of the VC. • <i>vc-id</i>—A unique value defined for the VC. The VC-ID connects the endpoints of the VC and must be the same on both ends of the VC. The range is from 1 to 4294967295.
Step 10	end	Return to privileged EXEC mode.
Step 11	show mpls l2transport vc	Verify the configuration.
Step 12	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no mpls l2transport route** *destination vc-id* or **no xconnect** *destination vc-id* interface command to delete the EoMPLS tunnel.

This example shows how to configure an EoMPLS tunnel between switch PE1's VLAN 3 interface and PE2's VLAN 4 interface.

PE1 has an IP address 10.0.0.1/32, and PE2 has IP address 20.0.0.1/32. Both PE routers are configured with an MPLS connection to the MPLS core. The VC ID is 123.

Enter these commands on the PE1 switch:

```
Switch(config)# interface loopback0
Switch(config-if)# ip address 10.10.10.10 255.255.255.255
Switch(config-if)# exit
Switch(config)# interface vlan 3
Switch(config-if)# mpls ip
Switch(config-if)# mpls l2transport route 20.0.0.1 123
```

Enter these commands on the PE2 switch:


```
Switch(config)# interface loopback0
Switch(config-if)# ip address 20.20.20.20 255.255.255.255
Switch(config-if)# exit
Switch(config)# interface vlan 4
Switch(config-if)# mpls ip
Switch(config-if)# mpls l2transport route 10.0.0.1 123
```

Packet Flow in an EoMPLS Network

Figure 30-5 is an example of packet flow in an EoMPLS network. A customer port on PE1 is configured for a per-port EoMPLS tunnel to a remote customer port on PE2. This allows the two physically separated customer switches (A and B) connected to these ports to appear as if they are directly connected on the same physical LAN.

The EoMPLS tunnel is configured with the IP address of Switch B and a VC ID that is associated with the remote customer port. PE1 establishes a tunnel LSP with PE2 by using a label advertised with LDP by Router A, which is connected to the ES port of PE1. PE1 then establishes a targeted LDP session to PE2 to advertise the VC label associated with the VC ID. When PE2 is configured with the EoMPLS tunnel, it also establishes a targeted LDP session to advertise the VC label it associated to the VC ID. This establishes an EoMPLS tunnel between the two ES ports on switch PE1 and switch PE2.

Figure 30-5 Sample EoMPLS Packet Flow



Assume that Host A is connected to the customer switch on VLAN 3 that has a trunk port connected to PE1 configured for 802.1Q tagging. Host A sends a packet to Host B, using the specific values of MAC addresses, labels, and VLANs shown in the figure. The customer switch tags the host packet and forwards it over the trunk port to PE1.

The tagged packet is received on the CE port that is configured for per-port EoMPLS tunneling. The PE1 switch examines the packet headers and looks at the tables stored in the switch to determine what to do with the packet. Because the port is configured for per-port EoMPLS tunneling, the switch does not remove any VLAN tags that are in the packet, but assigns the packet to an internal VLAN. Only the customer port and the ES port are configured with that internal VLAN, which makes the PE1 ES port the only possible destination for the packet.

The ES port encapsulates the packet header with the tunnel label and the VC label and forwards the packet to the next hop, in this case Router A, for it to be sent across the MPLS network.

The router receives the packet and forwards it over the MPLS network to the remote PE2 switch. PE2 removes the MPLS encapsulation and sends the packet out of the port associated with the VC label. Customer Switch B removes the final VLAN tag and forwards the packet to the remote host B.

VLAN-based EoMPLS packet flow is basically the same as port-based EoMPLS, except that the customer VLAN is used instead of an internal VLAN. The PE1 switch looks up the customer VLAN ID to determine that the packet is forwarded to the ES port, where the packets is again examined and encapsulated with the tunnel label and VC label based on the EoMPLS for that VLAN.

Configuring MPLS and EoMPLS QoS

Quality of service (QoS) in MPLS and EoMPLS enables network administrators to provide differentiated types of service across an MPLS network. Each packet can receive the particular kind of service specified by the packet QoS. To preserve QoS IP precedence bits, you must globally disable QoS.

After you enable QoS, you can preserve Differentiated Services Code Point (DSCP) or IP precedence bits by using a trusted configuration at the interface level. For more information, see the [“Configuring Ingress Classification by Using Port Trust States” section on page 26-42](#). However, the unpreserved bits are automatically overwritten by the value of the preserved bits. For example, if you preserve the DSCP bits, the IP precedence and CoS bits are overwritten with the value of the DSCP bits. You can also set MPLS and EoMPLS QoS priority by using 3 experimental bits in the MPLS label to determine the priority of packets.



Note

The switch supports only DSCP and IP precedence classification for MPLS and EoMPLS.

This section contains this information:

- [Understanding MPLS QoS, page 30-18](#)
- [Enabling MPLS and EoMPLS QoS, page 30-20](#)

Understanding MPLS QoS


Service in an MPLS network can be specified in different ways, for example, by using the IP precedence bit settings in IP packets. When you send IP packets from one site to another, the IP precedence field (the first 3 bits of the DSCP field in the header of an IP packet) specifies the QoS. Based on the IP precedence marking, the packet is given the desired treatment such as latency or bandwidth. If the network is an MPLS network, the IP precedence bits are copied into the MPLS EXP field at the edge of the network.

A service provider might also want to set a QoS value for an MPLS packet to a different value. Instead of overwriting the value in the IP precedence field that belongs to a customer, the service provider can set the MPLS experimental field. The IP header remains available for the customer's use, and the QoS of an IP packet is not changed as the packet travels through the MPLS network.

By choosing different values for the MPLS experimental field, you can mark packets based on their characteristics, such as rate or type, so that packets have the priority that they require during periods of congestion.

Figure 30-6 shows an MPLS network that connects two sites of an IP network that belongs to a customer.

Figure 30-6 MPLS Network Connecting Two Customer Sites



PE1 and PE2 are at the boundaries between the MPLS network and the IP network and are the ingress and egress PE devices. CE1 and CE2 are customer edge devices. P1 and P2 are service provider routers within the core of the service-provider network.

Packets arrive as IP packets at PE1, the ingress provider edge router, and PE1 sends the packets to the MPLS network as MPLS packets. Within the service-provider network, there is no IP precedence field for the queuing mechanism to look at because the packets are MPLS packets. The packets remain as MPLS packets until they arrive at PE2, the egress PE router. PE2 removes the label from each packet and forwards the packets as IP packets.

MPLS QoS enables service providers to classify packets according to their type, input interface, and other factors by setting (marking) each packet within the MPLS experimental field without changing the IP precedence or DSCP field. You can use the IP Precedence or DSCP bits to specify the QoS for an IP packet and use the MPLS experimental bits to specify the QoS for an MPLS packet. In an MPLS network, configure the MPLS experimental field value at PE1 (the ingress router) to set the QoS value in the MPLS packet.

It is important to assign the correct priority to a packet. The priority of a packet affects how the packet is treated during periods of congestion. For example, service providers have service-level agreements with customers that specify how much traffic the service provider has agreed to deliver. To comply with the agreement, the customer must not send more than the agreed-upon rate. Packets are considered to be in-rate or out-of-rate. If there is congestion in the network, out-of-rate packets might be dropped more aggressively.

Enabling MPLS and EoMPLS QoS

This section describes how to configure MPLS QoS on the ingress PE router. It includes these topics:

- [Default MPLS and EoMPLS QoS Configuration, page 30-20](#)
- [Setting the Priority of Packets with Experimental Bits, page 30-20](#)

For more information about QoS, see [Chapter 26, “Configuring QoS.”](#)

Default MPLS and EoMPLS QoS Configuration

QoS is disabled. Packets are not modified, and the CoS, DSCP, and IP precedence values in the packet are not changed. Traffic is switched in pass-through mode (packets are switched without any rewrites and classified as best effort without any policing).

The default behavior for the VLAN-tunneled EoMPLS or VFI packets is to relay the 802.1p bits into the EXP bits. The default value for the EXP bits for port-tunneled EoMPLS or VFI is zero.

When QoS is enabled with the **mls qos** global configuration command and all other QoS settings are at their defaults, traffic is classified as best effort (the DSCP value is set to 0) without any policing. No policy maps are configured.

**Note**

For MPLS and EoMPLS QoS, you can match only Layer 3 parameters (IP or DSCP values), not Layer 2 (CoS values).

Setting the Priority of Packets with Experimental Bits

MPLS and EoMPLS provide QoS on the ingress router by using 3 experimental bits in a label to determine the priority of packets. To support QoS between LERs, set the experimental bits in both the VC and tunnel labels. If you do not assign values to the experimental bits, the priority bits in the 802.1Q header tag control information field are written into the experimental bit fields.

The process includes these steps on the ingress router:

- Configure a class map to classify IP packets according to their DSCP or IP precedence classification.

**Note**

The switch supports only DSCP and IP precedence classification for MPLS and EoMPLS.

- Configure a policy map to mark MPLS packets (write their classification into the MPLS experimental field).
- Configure the input interface to attach the service policy.

Beginning in privileged EXEC mode, follow these steps to set the experimental bits for EoMPLS or MPLS QoS:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos	Enable QoS globally. QoS runs from the default settings described in Chapter 26, “Configuring QoS.”

	Command	Purpose
Step 3	class-map <i>class-map-name</i>	Specify the name of the traffic class, and enter class-map configuration mode.
Step 4	match { ip dscp <i>dscp-list</i> ip precedence <i>ip-precedence-list</i> }	Specify the matching criteria for 802.1Q packets. ip dscp <i>dscp-list</i>—list of up to eight IP DSCP values to match against incoming packets. The range is 0 to 63. ip precedence <i>ip-precedence-list</i>—list of up to eight IP-precedence values to match against incoming packets. Separate each value with a space. The range is 0 to 7. Note The switch does not support the cos and vlan keywords for MPLS and EoMPLS.
Step 5	exit	Return to global configuration mode.
Step 6	policy-map <i>policy-map-name</i>	Specify the name of the traffic policy to configure, and enter policy-map configuration mode.
Step 7	class <i>class-name</i>	Specify the name of the predefined traffic class configured with the class-map command, and enter policy-map class configuration mode.
Step 8	set mpls experimental <i>exp-number</i>	Specify the value to which the MPLS bits are set if the packets match the specified policy map. The range is 0 to 7.
Step 9	exit	Return to policy-map configuration mode.
Step 10	exit	Return to global configuration mode.
Step 11	interface <i>interface-id</i>	Enter the interface ID, and enter interface configuration mode. The interface should be the ES egress port of the ingress router.
Step 12	service-policy output <i>policy-map-name</i>	Attach the specified policy map to the output interface.
Step 13	end	Return to privileged EXEC mode.
Step 14	show policy-map [<i>policy-map-name</i> [<i>class class-map-name</i>]] show policy-map interface <i>interface-id</i>	Verify the configuration.
Step 15	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, use the **no class** *class-name* policy-map configuration command.

This example shows how to use class and policy maps to configure different experimental bit settings for DSCP and IP precedence for MPLS QoS.

```
Switch(config)# class-map match-all gold-class
Switch(config-cmap)# match ip dscp 1
Switch(config-cmap)# exit
Switch(config)# class-map match-all silver-class
Switch(config-cmap)# match ip precedence 2
Switch(config-cmap)# exit
```

```
Switch(config)# policy-map out-policy
Switch(config-pmap)# class gold-class
Switch(config-pmap-c)# set mpls experimental 5
Switch(config-pmap-c)# exit
Switch(config-pmap)# class silver-class
Switch(config-pmap-c)# set mpls experimental 4
Switch(config-pmap-c)# exit
```

```
Switch(config)# interface gigabitethernet1/1/1
Switch(config-if)# service-policy output out-policy
Switch(config-if)# end
```

Monitoring and Maintaining MPLS and EoMPLS

To clear MPLS counters or display MPLS and EoMPLS information, use the privileged EXEC commands in [Table 30-1](#).

Table 30-1 Commands for Displaying MPLS and EoMPLS Information

Command	Purpose
clear mpls counters	Clear MPLS forwarding counters.
show mpls forwarding-table	Display the contents of the MPLS label forwarding information base (LFIB).
show mpls interfaces	Display information about interfaces that have been configured for label switching.
show mpls ip binding	Display specified information about label bindings learned by LDP.
show mpls l2transport vc [detail] [summary]	Display detailed or summary information about the EoMPLS virtual connections that have been enabled to route Layer 2 packets on a provider-edge device.
show mpls l2transport vc [vc-id] [vc-id-min - vc-id-max]	Display information about the specified VC or range of VCs. The range is from 1 to 4294967295.
show mpls label range	Display the range of local labels available for use on packet interfaces.
show mpls ldp bindings	Display the contents of the label information base (LIB).
show mpls ldp discovery	Display the status of the LDP discovery process.
show mpls ldp neighbor	Display the status of LDP sessions.
show mpls ldp parameters	Display current LDP parameters.
show mpls prefix-map	Show the prefix map used to assign a QoS map to network prefixes that match a standard IP access list.
show mpls ldp backoff	Display information about the configured session setup backoff parameters and any potential LDP peers with which session setup attempts are being throttled.



Configuring IP Multicast Routing

This chapter describes how to configure IP multicast routing on the Catalyst 3750 Metro switch. IP multicasting is a more efficient way to use network resources, especially for bandwidth-intensive services such as audio and video. IP multicast routing enables a host (source) to send packets to a group of hosts (receivers) anywhere within the IP network by using a special form of IP address called the *IP multicast group address*. The sending host inserts the multicast group address into the IP destination address field of the packet, and IP multicast routers and multilayer switches forward incoming IP multicast packets out all interfaces that lead to members of the multicast group. Any host, regardless of whether it is a member of a group, can send to a group. However, only the members of a group receive the message.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

This chapter consists of these sections:

- [Understanding Cisco's Implementation of IP Multicast Routing, page 31-2](#)
- [Configuring IP Multicast Routing, page 31-8](#)
- [Configuring Advanced PIM Features, page 31-22](#)
- [Configuring Optional IGMP Features, page 31-26](#)
- [Configuring Optional Multicast Routing Features, page 31-31](#)
- [Configuring Basic DVMRP Interoperability Features, page 31-36](#)
- [Configuring Advanced DVMRP Interoperability Features, page 31-41](#)
- [Monitoring and Maintaining IP Multicast Routing, page 31-49](#)

For information on configuring the Multicast Source Discovery Protocol (MSDP), see [Chapter 32, "Configuring MSDP."](#)


Understanding Cisco's Implementation of IP Multicast Routing

The Cisco IOS software supports these protocols to implement IP multicast routing:

- Internet Group Management Protocol (IGMP) is used among hosts on a LAN and the routers (and multilayer switches) on that LAN to track the multicast groups of which hosts are members.
- Protocol-Independent Multicast (PIM) protocol is used among routers and multilayer switches to track which multicast packets to forward to each other and to their directly connected LANs.
- Distance Vector Multicast Routing Protocol (DVMRP) is used on the multicast backbone of the Internet (MBONE). The software supports PIM-to-DVMRP interaction.
- Cisco Group Management Protocol (CGMP) is used on Cisco routers and multilayer switches connected to Layer 2 Catalyst switches to perform tasks similar to those performed by IGMP.

Figure 31-1 shows where these protocols operate within the IP multicast environment.

Figure 31-1 IP Multicast Routing Protocols



Understanding IGMP

To participate in IP multicasting, multicast hosts, routers, and multilayer switches must have the IGMP operating. This protocol defines the querier and host roles:

- A querier is a network device that sends query messages to discover which network devices are members of a given multicast group.
- A host is a receiver that sends report messages (in response to query messages) to inform a querier of a host membership.

A set of queriers and hosts that receive multicast data streams from the same source is called a multicast group. Queriers and hosts use IGMP messages to join and leave multicast groups.

Any host, regardless of whether it is a member of a group, can send to a group. However, only the members of a group receive the message. Membership in a multicast group is dynamic; hosts can join and leave at any time. There is no restriction on the location or number of members in a multicast group. A host can be a member of more than one multicast group at a time. How active a multicast group is and

what members it has can vary from group to group and from time to time. A multicast group can be active for a long time, or it can be very short-lived. Membership in a group can constantly change. A group that has members can have no activity.

IP multicast traffic uses group addresses, which are class D addresses. The high-order bits of a Class D address are 1110. Therefore, host group addresses can be in the range 224.0.0.0 through 239.255.255.255. Multicast addresses in the range 224.0.0.0 to 240.0.0.255 are reserved for use by routing protocols and other network control traffic. The address 224.0.0.0 is guaranteed not to be assigned to any group.

IGMP packets are sent using these IP multicast group addresses:

- IGMP general queries are destined to the address 224.0.0.1 (all systems on a subnet).
- IGMP group-specific queries are destined to the group IP address for which the switch is querying.
- IGMP group membership reports are destined to the group IP address for which the switch is reporting.
- IGMP Version 2 (IGMPv2) leave messages are destined to the address 224.0.0.2 (all-multicast-routers on a subnet). In some old host IP stacks, leave messages might be destined to the group IP address rather than to the all-routers address.

IGMP Version 1

IGMP Version 1 (IGMPv1) primarily uses a query-response model that enables the multicast router and multilayer switch to find which multicast groups are active (have one or more hosts interested in a multicast group) on the local subnet. IGMPv1 has other processes that enable a host to join and leave a multicast group. For more information, refer to RFC 1112.

IGMP Version 2

IGMPv2 extends IGMP functionality by providing such features as the IGMP leave process to reduce leave latency, group-specific queries, and an explicit maximum query response time. IGMPv2 also adds the capability for routers to elect the IGMP querier without depending on the multicast protocol to perform this task. For more information, refer to RFC 2236.

Understanding PIM

PIM is called *protocol-independent*: regardless of the unicast routing protocols used to populate the unicast routing table, PIM uses this information to perform multicast forwarding instead of maintaining a separate multicast routing table.

PIM is defined in RFC 2362, *Protocol-Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification*. PIM is defined in these Internet Engineering Task Force (IETF) Internet drafts:

- *Protocol Independent Multicast (PIM): Motivation and Architecture*
- *Protocol Independent Multicast (PIM), Dense Mode Protocol Specification*
- *Protocol Independent Multicast (PIM), Sparse Mode Protocol Specification*
- *draft-ietf-idmr-igmp-v2-06.txt, Internet Group Management Protocol, Version 2*
- *draft-ietf-pim-v2-dm-03.txt, PIM Version 2 Dense Mode*

PIM Versions

PIMv2 includes these improvements over PIMv1:

- A single, active rendezvous point (RP) exists per multicast group, with multiple backup RPs. This single RP compares to multiple active RPs for the same group in PIMv1.
- A bootstrap router (BSR) provides a fault-tolerant, automated RP discovery and distribution mechanism that enables routers and multilayer switches to dynamically learn the group-to-RP mappings.
- Sparse mode and dense mode are properties of a group, as opposed to an interface. We strongly recommend sparse-dense mode, as opposed to either sparse mode or dense mode only.
- PIM join and prune messages have more flexible encoding for multiple address families.
- A more flexible hello packet format replaces the query packet to encode current and future capability options.
- Register messages to an RP specify whether they are sent by a border router or a designated router.
- PIM packets are no longer inside IGMP packets; they are standalone packets.

PIM Modes

PIM can operate in dense mode (DM), sparse mode (SM), or in sparse-dense mode (PIM DM-SM), which handles both sparse groups and dense groups at the same time.

PIM DM

PIM DM builds source-based multicast distribution trees. In dense mode, a PIM DM router or multilayer switch assumes that all other routers or multilayer switches forward multicast packets for a group. If a PIM DM device receives a multicast packet and has no directly connected members or PIM neighbors present, a prune message is sent back to the source to stop unwanted multicast traffic. Subsequent multicast packets are not flooded to this router or switch on this pruned branch because branches without receivers are pruned from the distribution tree, leaving only branches that contain receivers.

When a new receiver on a previously pruned branch of the tree joins a multicast group, the PIM DM device detects the new receiver and immediately sends a graft message up the distribution tree toward the source. When the upstream PIM DM device receives the graft message, it immediately puts the interface on which the graft was received into the forwarding state so that the multicast traffic begins flowing to the receiver.

PIM SM

PIM SM uses shared trees and shortest-path-trees (SPTs) to distribute multicast traffic to multicast receivers in the network. In PIM SM, a router or multilayer switch assumes that other routers or switches do not forward multicast packets for a group, unless there is an explicit request for the traffic (join message). When a host joins a multicast group using IGMP, its directly connected PIM SM device sends PIM join messages toward the root, also known as the RP. This join message travels router-by-router toward the root, constructing a branch of the shared tree as it goes.

The RP keeps track of multicast receivers. It also registers sources through register messages received from the source's first-hop router (*designated router* [DR]) to complete the shared tree path from the source to the receiver. When using a shared tree, sources must send their traffic to the RP so that the traffic reaches all receivers.

Prune messages are sent up the distribution tree to prune multicast group traffic. This action permits branches of the shared tree or SPT that were created with explicit join messages to be torn down when they are no longer needed.

Auto-RP

This proprietary feature eliminates the need to manually configure the RP information in every router and multilayer switch in the network. For Auto-RP to work, you configure a Cisco router or multilayer switch as the mapping agent. It uses IP multicast to learn which routers or switches in the network are possible candidate RPs to receive candidate RP announcements. Candidate RPs periodically send multicast RP-announce messages to a particular group or group range to announce their availability.

Mapping agents listen to these candidate RP announcements and use the information to create entries in their Group-to-RP mapping caches. Only one mapping cache entry is created for any Group-to-RP range received, even if multiple candidate RPs are sending RP announcements for the same range. As the RP-announce messages arrive, the mapping agent selects the router or switch with the highest IP address as the active RP and stores this RP address in the Group-to-RP mapping cache.

Mapping agents periodically multicast the contents of their Group-to-RP mapping cache. Thus, all routers and switches automatically discover which RP to use for the groups they support. If a router or switch fails to receive RP-discovery messages and the Group-to-RP mapping information expires, it switches to a statically configured RP that was defined with the **ip pim rp-address** global configuration command. If no statically configured RP exists, the router or switch changes the group to dense-mode operation.

Multiple RPs serve different group ranges or serve as hot backups of each other.

Bootstrap Router

PIMv2 BSR is another method to distribute group-to-RP mapping information to all PIM routers and multilayer switches in the network. It eliminates the need to manually configure RP information in every router and switch in the network. However, instead of using IP multicast to distribute group-to-RP mapping information, BSR uses hop-by-hop flooding of special BSR messages to distribute the mapping information.

The BSR is elected from a set of candidate routers and switches in the domain that have been configured to function as BSRs. The election mechanism is similar to the root-bridge election mechanism used in bridged LANs. The BSR election is based on the BSR priority of the device contained in the BSR messages that are sent hop-by-hop through the network. Each BSR device examines the message and forwards out all interfaces only the message that has either a higher BSR priority than its BSR priority or the same BSR priority, but with a higher BSR IP address. Using this method, the BSR is elected.

The elected BSR sends BSR messages with a TTL of 1. Neighboring PIMv2 routers or multilayer switches receive the BSR message and multicast it out all other interfaces (except the one on which it was received) with a TTL of 1. In this way, BSR messages travel hop-by-hop throughout the PIM domain. Because BSR messages contain the IP address of the current BSR, the flooding mechanism enables candidate RPs to automatically learn which device is the elected BSR.

Candidate RPs send candidate RP advertisements showing the group range for which they are responsible to the BSR, which stores this information in its local candidate-RP cache. The BSR periodically advertises the contents of this cache in BSR messages to all other PIM devices in the domain. These messages travel hop-by-hop through the network to all routers and switches, which store the RP information in the BSR message in their local RP cache. The routers and switches select the same RP for a given group because they all use a common RP hashing algorithm.

Multicast Forwarding and Reverse Path Check

With unicast routing, routers and multilayer switches forward traffic through the network along a single path from the source to the destination host whose IP address appears in the destination address field of the IP packet. Each router and switch along the way makes a unicast forwarding decision, using the destination IP address in the packet, by looking up the destination address in the unicast routing table and forwarding the packet through the specified interface to the next hop toward the destination.

With multicasting, the source is sending traffic to an arbitrary group of hosts represented by a multicast group address in the destination address field of the IP packet. To decide whether to forward or drop an incoming multicast packet, the router or multilayer switch uses a reverse path forwarding (RPF) check on the packet as follows and shown in [Figure 31-2](#):

1. The router or multilayer switch examines the source address of the arriving multicast packet to decide whether the packet arrived on an interface that is on the reverse path back to the source.
2. If the packet arrives on the interface leading back to the source, the RPF check is successful and the packet is forwarded to all interfaces in the outgoing interface list (which might not be all interfaces on the router).
3. If the RPF check fails, the packet is discarded.

Some multicast routing protocols, such as DVMRP, maintain a separate multicast routing table and use it for the RPF check. However, PIM uses the unicast routing table to perform the RPF check.

[Figure 31-2](#) shows port 0/2 receiving a multicast packet from source 151.10.3.21. [Table 31-1](#) shows that the port on the reverse path to the source is port 0/1, not port 0/2. Because the RPF check fails, the multilayer switch discards the packet. Another multicast packet from source 151.10.3.21 is received on port 0/1, and the routing table shows this port is on the reverse path to the source. Because the RPF check passes, the switch forwards the packet to all ports in the outgoing port list.

Figure 31-2 RPF Check




Table 31-1 Routing Table Example for an RPF Check

Network	Port
151.10.0.0/16	Gigabit Ethernet 1/0/1
198.14.32.0/32	Gigabit Ethernet 1/1/1
204.1.16.0/24	Gigabit Ethernet 1/1/2

PIM uses both source trees and RP-rooted shared trees to forward datagrams (described in the “[PIM DM](#)” section on page 31-4 and the “[PIM SM](#)” section on page 31-4). The RPF check is performed differently for each:

- If a PIM router or multilayer switch has a source-tree state (that is, an (S,G) entry is present in the multicast routing table), it performs the RPF check against the IP address of the source of the multicast packet.
- If a PIM router or multilayer switch has a shared-tree state (and no explicit source-tree state), it performs the RPF check on the RP address (which is known when members join the group).

Sparse-mode PIM uses the RPF lookup function to decide where it needs to send joins and prunes:

- (S,G) joins (which are source-tree states) are sent toward the source.
- (*,G) joins (which are shared-tree states) are sent toward the RP.

DVMRP and dense-mode PIM use only source trees and use RPF as previously described.

Understanding DVMRP

DVMRP is implemented in the equipment of many vendors and is based on the public-domain mrouterd program. This protocol has been deployed in the MBONE and in other intradomain multicast networks.

Cisco routers and multilayer switches run PIM and can forward multicast packets to and receive from a DVMRP neighbor. It is also possible to propagate DVMRP routes into and through a PIM cloud. The software propagates DVMRP routes and builds a separate database for these routes on each router and multilayer switch, but PIM uses this routing information to make the packet-forwarding decision. The software does not implement the complete DVMRP. However, it supports dynamic discovery of DVMRP routers and can interoperate with them over traditional media (such as Ethernet and FDDI) or over DVMRP-specific tunnels.

DVMRP neighbors build a route table by periodically exchanging source network routing information in route-report messages. The routing information stored in the DVMRP routing table is separate from the unicast routing table and is used to build a source distribution tree and to perform multicast forward using RPF.

DVMRP is a dense-mode protocol and builds a parent-child database using a constrained multicast model to build a forwarding tree rooted at the source of the multicast packets. Multicast packets are initially flooded down this source tree. If redundant paths are on the source tree, packets are not forwarded along those paths. Forwarding occurs until prune messages are received on those parent-child links, which further constrain the broadcast of multicast packets.

Understanding CGMP

This software release provides CGMP-server support on your switch; no client-side functionality is provided. The switch serves as a CGMP server for devices that do not support IGMP snooping but have CGMP-client functionality.

CGMP is a protocol used on Cisco routers and multilayer switches connected to Layer 2 Catalyst switches to perform tasks similar to those performed by IGMP. CGMP permits Layer 2 group membership information to be communicated from the CGMP server to the switch. The switch can then learn on which interfaces multicast members reside instead of flooding multicast traffic to all switch interfaces. (IGMP snooping is another method to constrain the flooding of multicast packets. For more information, see [Chapter 17, “Configuring IGMP Snooping and MVR.”](#))

CGMP is necessary because the Layer 2 switch cannot distinguish between IP multicast data packets and IGMP report messages, which are both at the MAC-level and are addressed to the same group address.

Configuring IP Multicast Routing

These sections describe how to configure IP multicast routing:

- [Default Multicast Routing Configuration, page 31-8](#)
- [Multicast Routing Configuration Guidelines, page 31-8](#)
- [Configuring Basic Multicast Routing, page 31-10](#) (required)
- [Configuring a Rendezvous Point, page 31-11](#) (required if the interface is in sparse-dense mode, and you want to treat the group as a sparse group)
- [Using Auto-RP and a BSR, page 31-21](#) (required for non-Cisco PIMv2 devices to interoperate with Cisco PIM v1 devices))
- [Monitoring the RP Mapping Information, page 31-22](#) (optional)
- [Troubleshooting PIMv1 and PIMv2 Interoperability Problems, page 31-22](#) (optional)

Default Multicast Routing Configuration

[Table 31-2](#) shows the default multicast routing configuration.

Table 31-2 Default Multicast Routing Configuration

Feature	Default Setting
Multicast routing	Disabled on all interfaces.
PIM version	Version 2.
PIM mode	No mode is defined.
PIM RP address	None configured.
PIM domain border	Disabled.
PIM multicast boundary	None.
Candidate BSRs	Disabled.
Candidate RPs	Disabled.
Shortest-path tree threshold rate	0 kbps.
PIM router query message interval	30 seconds.

Multicast Routing Configuration Guidelines

To avoid misconfiguring multicast routing on your switch, review the information in these sections:

- [PIMv1 and PIMv2 Interoperability, page 31-9](#)
- [Auto-RP and BSR Configuration Guidelines, page 31-9](#)

PIMv1 and PIMv2 Interoperability

The Cisco PIMv2 implementation provides interoperability and transition between Version 1 and Version 2, although there might be some minor problems.

You can upgrade to PIMv2 incrementally. PIM Versions 1 and 2 can be configured on different routers and multilayer switches within one network. Internally, all routers and multilayer switches on a shared media network must run the same PIM version. Therefore, if a PIMv2 device detects a PIMv1 device, the Version 2 device downgrades itself to Version 1 until all Version 1 devices have been shut down or upgraded.

PIMv2 uses the BSR to discover and announce RP-set information for each group prefix to all the routers and multilayer switches in a PIM domain. PIMv1, together with the Auto-RP feature, can perform the same tasks as the PIMv2 BSR. However, Auto-RP is a standalone protocol, separate from PIMv1, and is a proprietary Cisco protocol. PIMv2 is a standards track protocol in the IETF. We recommend that you use PIMv2. The BSR mechanism interoperates with Auto-RP on Cisco routers and multilayer switches. For more information, see the [“Auto-RP and BSR Configuration Guidelines” section on page 31-9](#).

When PIMv2 devices interoperate with PIMv1 devices, Auto-RP should have already been deployed. A PIMv2 BSR that is also an Auto-RP mapping agent automatically advertises the RP elected by Auto-RP. That is, Auto-RP sets its single RP on every router or multilayer switch in the group. Not all routers and switches in the domain use the PIMv2 hash function to select multiple RPs.

Dense-mode groups in a mixed PIMv1 and PIMv2 region need no special configuration; they automatically interoperate.

Sparse-mode groups in a mixed PIMv1 and PIMv2 region are possible because the Auto-RP feature in PIMv1 interoperates with the PIMv2 RP feature. Although all PIMv2 devices can also use PIMv1, we recommend that the RPs be upgraded to PIMv2. To ease the transition to PIMv2, we have these recommendations:

- Use Auto-RP throughout the region.
- Configure sparse-dense mode throughout the region.

If Auto-RP is not already configured in the PIMv1 regions, configure Auto-RP. For more information, see the [“Configuring Auto-RP” section on page 31-13](#).

Auto-RP and BSR Configuration Guidelines

There are two approaches to using PIMv2. You can use Version 2 exclusively in your network or migrate to Version 2 by employing a mixed PIM version environment.

- If your network is all Cisco routers and multilayer switches, you can use either Auto-RP or BSR.
- If you have non-Cisco routers in your network, you must use BSR.
- If you have Cisco PIMv1 and PIMv2 routers and multilayer switches and non-Cisco routers, you must use both Auto-RP and BSR. If your network includes routers from other vendors, configure the Auto-RP mapping agent and the BSR on a Cisco PIMv2 device. Ensure that no PIMv1 device is located in the path between the BSR and a non-Cisco PIMv2 device.
- Because bootstrap messages are sent hop-by-hop, a PIMv1 device prevents these messages from reaching all routers and multilayer switches in your network. Therefore, if your network has a PIMv1 device in it and only Cisco routers and multilayer switches, it is best to use Auto-RP.

- If you have a network that includes non-Cisco routers, configure the Auto-RP mapping agent and the BSR on a Cisco PIMv2 router or multilayer switch. Ensure that no PIMv1 device is on the path between the BSR and a non-Cisco PIMv2 router.
- If you have non-Cisco PIMv2 routers that need to interoperate with Cisco PIMv1 routers and multilayer switches, both Auto-RP and a BSR are required. We recommend that a Cisco PIMv2 device be both the Auto-RP mapping agent and the BSR. For more information, see the [“Using Auto-RP and a BSR” section on page 31-21](#).

Configuring Basic Multicast Routing

You must enable IP multicast routing and configure the PIM version and the PIM mode. Then the software can forward multicast packets, and the switch can populate its multicast routing table.

You can configure an interface to be in PIM dense mode, sparse mode, or sparse-dense mode. The switch populates its multicast routing table and forwards multicast packets it receives from its directly connected LANs according to the mode setting. You must enable PIM in one of these modes for an interface to perform IP multicast routing. Enabling PIM on an interface also enables IGMP operation on that interface.

In populating the multicast routing table, dense-mode interfaces are always added to the table. Sparse-mode interfaces are added to the table only when periodic join messages are received from downstream devices or when there is a directly connected member on the interface. When forwarding from a LAN, sparse-mode operation occurs if there is an RP known for the group. If so, the packets are encapsulated and sent toward the RP. When no RP is known, the packet is flooded in a dense-mode fashion. If the multicast traffic from a specific source is sufficient, the receiver’s first-hop router might send join messages toward the source to build a source-based distribution tree.

By default, multicast routing is disabled, and there is no default mode setting. This procedure is required. Beginning in privileged EXEC mode, follow these steps to enable IP multicasting, to configure a PIM version, and to configure a PIM mode. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip multicast-routing distributed	Enable IP multicast distributed switching.
Step 3	interface <i>interface-id</i>	Specify the Layer 3 interface on which you want to enable multicast routing, and enter interface configuration mode. The specified interface must be one of the following: • A routed port: a physical port that has been configured as a Layer 3 port by entering the no switchport interface configuration command. • An SVI: a VLAN interface created by using the interface vlan <i>vlan-id</i> global configuration command. These interfaces must have IP addresses assigned to them. For more information, see the “Configuring Layer 3 Interfaces” section on page 9-18.

	Command	Purpose
Step 4	ip pim version [1 2]	Configure the PIM version on the interface. By default, Version 2 is enabled and is the recommended setting. An interface in PIMv2 mode automatically downgrades to PIMv1 mode if that interface has a PIMv1 neighbor. The interface returns to Version 2 mode after all Version 1 neighbors are shut down or upgraded. For more information, see the “PIMv1 and PIMv2 Interoperability” section on page 31-9.
Step 5	ip pim {dense-mode sparse-mode sparse-dense-mode}	Enable a PIM mode on the interface. By default, no mode is configured. The keywords have these meanings: • dense-mode—Enables dense mode of operation. • sparse-mode—Enables sparse mode of operation. If you configure sparse-mode, you must also configure an RP. For more information, see the “Configuring a Rendezvous Point” section on page 31-11. • sparse-dense-mode—Causes the interface to be treated in the mode in which the group belongs. Sparse-dense-mode is the recommended setting. Note After you enable a PIM mode on the interface, the ip mroute-cache distributed interface configuration command is automatically entered for the interface and appears in the running configuration.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable multicasting, use the **no ip multicast-routing distributed** global configuration command. To return to the default PIM version, use the **no ip pim version** interface configuration command. To disable PIM on an interface, use the **no ip pim** interface configuration command.

Configuring a Rendezvous Point

You must have an RP if the interface is in sparse-dense mode and if you want to treat the group as a sparse group. You can use several methods, as described in these sections:

- [Manually Assigning an RP to Multicast Groups, page 31-12](#)
- [Configuring Auto-RP, page 31-13](#) (a standalone, Cisco-proprietary protocol separate from PIMv1)
- [Configuring PIMv2 BSR, page 31-17](#) (a standards track protocol in the IETF)

You can use Auto-RP, BSR, or a combination of both, depending on the PIM version you are running and the types of routers in your network. For more information, see the [“PIMv1 and PIMv2 Interoperability” section on page 31-9](#) and the [“Auto-RP and BSR Configuration Guidelines” section on page 31-9](#).

Manually Assigning an RP to Multicast Groups

This section explains how to manually configure an RP. If the RP for a group is learned through a dynamic mechanism (such as Auto-RP or BSR), you need not perform this task for that RP.

Senders of multicast traffic announce their existence through register messages received from the source's first-hop router (designated router) and forwarded to the RP. Receivers of multicast packets use RPs to join a multicast group by using explicit join messages. RPs are not members of the multicast group; rather, they serve as a *meeting place* for multicast sources and group members.

You can configure a single RP for multiple groups defined by an access list. If there is no RP configured for a group, the multilayer switch treats the group as dense and uses the dense-mode PIM techniques.

Beginning in privileged EXEC mode, follow these steps to manually configure the address of the RP. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip pim rp-address <i>ip-address</i> [<i>access-list-number</i>] [override]	Configure the address of a PIM RP. By default, no PIM RP address is configured. You must configure the IP address of RPs on all routers and multilayer switches (including the RP). If there is no RP configured for a group, the switch treats the group as dense, using the dense-mode PIM techniques. A PIM device can be an RP for more than one group. Only one RP address can be used at a time within a PIM domain. The access-list conditions specify for which groups the device is an RP. - For <i>ip-address</i>, enter the unicast address of the RP in dotted-decimal notation. (Optional) For <i>access-list-number</i>, enter an IP standard access list number from 1 to 99. If no access list is configured, the RP is used for all groups. (Optional) The override keyword means that if there is a conflict between the RP configured with this command and one learned by Auto-RP or BSR, the RP configured with this command prevails.
Step 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. - 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 multicast group address for which the RP should be used. (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.
Step 4	end	Return to privileged EXEC mode.

	Command	Purpose
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove an RP address, use the **no ip pim rp-address** *ip-address* [*access-list-number*] [**override**] global configuration command.

This example shows how to configure the address of the RP to 147.106.6.22 for multicast group 225.2.2.2 only:

```
Switch(config)# access-list 1 permit 225.2.2.2 0.0.0.0
Switch(config)# ip pim rp-address 147.106.6.22 1
```

Configuring Auto-RP

Auto-RP uses IP multicast to automate the distribution of group-to-RP mappings to all Cisco routers and multilayer switches in a PIM network. It has these benefits:

- It is easy to use multiple RPs within a network to serve different group ranges.
- It provides load splitting among different RPs and arrangement of RPs according to the location of group participants.
- It avoids inconsistent, manual RP configurations on every router and multilayer switch in a PIM network, which can cause connectivity problems.



Note

If you configure PIM in sparse mode or sparse-dense mode and do not configure Auto-RP, you must manually configure an RP as described in the [“Manually Assigning an RP to Multicast Groups”](#) section on page 31-12.



Note

If routed ports are configured in sparse mode, Auto-RP can still be used if all devices are configured with a manual RP address for the Auto-RP groups.

These sections describe how to configure Auto-RP:

- [Setting up Auto-RP in a New Internetwork](#), page 31-13 (optional)
- [Adding Auto-RP to an Existing Sparse-Mode Cloud](#), page 31-14 (optional)
- [Preventing Join Messages to False RPs](#), page 31-15 (optional)
- [Filtering Incoming RP Announcement Messages](#), page 31-16 (optional)

For overview information, see the [“Auto-RP”](#) section on page 31-5.

Setting up Auto-RP in a New Internetwork

If you are setting up Auto-RP in a new internetwork, you do not need a default RP because you configure all the interfaces for sparse-dense mode. Follow the process described in the next section [“Adding Auto-RP to an Existing Sparse-Mode Cloud”](#) section on page 31-14. However, skip Step 3 to configure a PIM router as the RP for the local group.

Adding Auto-RP to an Existing Sparse-Mode Cloud

This section contains some suggestions for the initial deployment of Auto-RP into an existing sparse-mode cloud to minimize disruption of the existing multicast infrastructure.

Beginning in privileged EXEC mode, follow these steps to deploy Auto-RP in an existing sparse-mode cloud. This procedure is optional.

	Command	Purpose
Step 1	show running-config	Verify that a default RP is already configured on all PIM devices and the RP in the sparse-mode network. It was previously configured with the ip pim rp-address global configuration command. This step is not required for sparse-dense-mode environments. The selected RP should have good connectivity and be available across the network. Use this RP for the global groups (for example 224.x.x.x and other global groups). Do not reconfigure the group address range that this RP serves. RPs dynamically discovered through Auto-RP take precedence over statically configured RPs. Assume that it is desirable to use a second RP for the local groups.
Step 2	configure terminal	Enter global configuration mode.
Step 3	ip pim send-rp-announce <i>interface-id</i> scope <i>ttl</i> group-list <i>access-list-number</i> interval <i>seconds</i>	Configure another PIM device to be the candidate RP for local groups. - For <i>interface-id</i>, enter the interface type and number that identifies the RP address. Valid interfaces include physical ports, port channels, and VLANs. - For scope <i>ttl</i>, specify the time-to-live value in hops. Enter a hop count that is high enough so that the RP-announce messages reach all mapping agents in the network. There is no default setting. The range is 1 to 255. - For group-list <i>access-list-number</i>, enter an IP standard access list number from 1 to 99. If no access list is configured, the RP is used for all groups. - For interval <i>seconds</i>, specify how often the announcement messages must be sent. The default is 60 seconds. The range is 1 to 16383.
Step 4	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. - For <i>access-list-number</i>, enter the access list number specified in Step 3. - 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 multicast group address range for which the RP should be used. (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.

	Command	Purpose
Step 5	ip pim send-rp-discovery scope ttl	Find a switch whose connectivity is not likely to be interrupted, and assign it the role of RP-mapping agent. For scope ttl , specify the time-to-live value in hops to limit the RP discovery packets. All devices within the hop count from the source device receive the Auto-RP discovery messages. These messages tell other devices which group-to-RP mapping to use to avoid conflicts (such as overlapping group-to-RP ranges). There is no default setting. The range is 1 to 255.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.
	show ip pim rp mapping	Display active RPs that are cached with associated multicast routing entries.
	show ip pim rp	Display the information cached in the routing table.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the PIM device configured as the candidate RP, use the **no ip pim send-rp-announce interface-id** global configuration command. To remove the switch as the RP-mapping agent, use the **no ip pim send-rp-discovery** global configuration command.

This example shows how to send RP announcements out all PIM-enabled ports for a maximum of 31 hops. The IP address of port 1 is the RP. Access list 5 describes the group for which this switch serves as RP:

```
Switch(config)# ip pim send-rp-announce gigabitethernet1/0/1 scope 31 group-list 5
Switch(config)# access-list 5 permit 224.0.0.0 15.255.255.255
```

Preventing Join Messages to False RPs

Find whether the **ip pim accept-rp** command was previously configured throughout the network by using the **show running-config** privileged EXEC command. If the **ip pim accept-rp** command is not configured on any device, this problem can be addressed later. In those routers or multilayer switches already configured with the **ip pim accept-rp** command, you must enter the command again to accept the newly advertised RP.

To accept all RPs advertised with Auto-RP and reject all other RPs by default, use the **ip pim accept-rp auto-rp** global configuration command. This procedure is optional.

If all ports are in sparse mode, use a default-configured RP to support the two well-known groups 224.0.1.39 and 224.0.1.40. Auto-RP uses these two well-known groups to collect and distribute RP-mapping information. When this is the case and the **ip pim accept-rp auto-rp** command is configured, another **ip pim accept-rp** command accepting the RP must be configured as follows:

```
Switch(config)# ip pim accept-rp 172.10.20.1 1
Switch(config)# access-list 1 permit 224.0.1.39
Switch(config)# access-list 1 permit 224.0.1.40
```

Filtering Incoming RP Announcement Messages

You can add configuration commands to the mapping agents to prevent a maliciously configured router from masquerading as a candidate RP and causing problems.

Beginning in privileged EXEC mode, follow these steps to filter incoming RP announcement messages. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip pim rp-announce-filter rp-list access-list-number group-list access-list-number	Filter incoming RP announcement messages. Enter this command on each mapping agent in the network. Without this command, all incoming RP-announce messages are accepted by default. For rp-list access-list-number, configure an access list of candidate RP addresses that, if permitted, is accepted for the group ranges supplied in the group-list access-list-number variable. If this variable is omitted, the filter applies to all multicast groups. If more than one mapping agent is used, the filters must be consistent across all mapping agents to ensure that no conflicts occur in the Group-to-RP mapping information.
Step 3	access-list access-list-number {deny permit} source [source-wildcard]	Create a standard access list, repeating the command as many times as necessary. - 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. - Create an access list that specifies from which routers and multilayer switches the mapping agent accepts candidate RP announcements (rp-list ACL). - Create an access list that specifies the range of multicast groups from which to accept or deny (group-list ACL). - For <i>source</i>, enter the multicast group address range for which the RP should be used. (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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a filter on incoming RP announcement messages, use the **no ip pim rp-announce-filter rp-list access-list-number [group-list access-list-number]** global configuration command.

This example shows a sample configuration on an Auto-RP mapping agent that is used to prevent candidate RP announcements from being accepted from unauthorized candidate RPs:

```
Switch(config)# ip pim rp-announce-filter rp-list 10 group-list 20
Switch(config)# access-list 10 permit host 172.16.5.1
Switch(config)# access-list 10 permit host 172.16.2.1
Switch(config)# access-list 20 deny 239.0.0.0 0.0.255.255
Switch(config)# access-list 20 permit 224.0.0.0 15.255.255.255
```

In this example, the mapping agent accepts candidate RP announcements from only two devices, 172.16.5.1 and 172.16.2.1. The mapping agent accepts candidate RP announcements from these two devices only for multicast groups that fall in the group range of 224.0.0.0 to 239.255.255.255. The mapping agent does not accept candidate RP announcements from any other devices in the network. Furthermore, the mapping agent does not accept candidate RP announcements from 172.16.5.1 or 172.16.2.1 if the announcements are for any groups in the 239.0.0.0 through 239.255.255.255 range. This range is the administratively scoped address range.

Configuring PIMv2 BSR

These sections describe how to set up BSR in your PIMv2 network:

- [Defining the PIM Domain Border, page 31-17](#) (optional)
- [Defining the IP Multicast Boundary, page 31-18](#) (optional)
- [Configuring Candidate BSRs, page 31-19](#) (optional)
- [Configuring Candidate RPs, page 31-20](#) (optional)

For overview information, see the “Bootstrap Router” section on page 31-5.

Defining the PIM Domain Border


As IP multicast becomes more widespread, the chance of one PIMv2 domain bordering another PIMv2 domain is increasing. Because these two domains probably do not share the same set of RPs, BSR, candidate RPs, and candidate BSRs, you need to constrain PIMv2 BSR messages from flowing into or out of the domain. Allowing these messages to leak across the domain borders could adversely affect the normal BSR election mechanism and elect a single BSR across all bordering domains and co-mingle candidate RP advertisements, resulting in the election of RPs in the wrong domain.

Beginning in privileged EXEC mode, follow these steps to define the PIM domain border. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to be configured, and enter interface configuration mode.
Step 3	ip pim bsr-border	Define a PIM bootstrap message boundary for the PIM domain. Enter this command on each interface that connects to other bordering PIM domains. This command instructs the switch to neither send or receive PIMv2 BSR messages on this interface as shown in Figure 31-3 .
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the PIM border, use the **no ip pim bsr-border** interface configuration command.

Figure 31-3 Constraining PIMv2 BSR Messages



101243

Defining the IP Multicast Boundary

You define a multicast boundary to prevent Auto-RP messages from entering the PIM domain. You create an access list to deny packets destined for 224.0.1.39 and 224.0.1.40, which carry Auto-RP information.

Beginning in privileged EXEC mode, follow these steps to define a multicast boundary. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	access-list access-list-number deny source [source-wildcard]	Create a standard access list, repeating the command as many times as necessary. - For <i>access-list-number</i>, the range is 1 to 99. - The deny keyword denies access if the conditions are matched. - For <i>source</i>, enter multicast addresses 224.0.1.39 and 224.0.1.40, which carry Auto-RP information. (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.
Step 3	interface interface-id	Specify the interface to be configured, and enter interface configuration mode.
Step 4	ip multicast boundary access-list-number	Configure the boundary, specifying the access list you created in Step 2.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the boundary, use the **no ip multicast boundary** interface configuration command.

This example shows a portion of an IP multicast boundary configuration that denies Auto-RP information:

```
Switch(config)# access-list 1 deny 224.0.1.39
Switch(config)# access-list 1 deny 224.0.1.40
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# ip multicast boundary 1
```

Configuring Candidate BSRs

You can configure one or more candidate BSRs. The devices serving as candidate BSRs should have good connectivity to other devices and be in the backbone portion of the network.

Beginning in privileged EXEC mode, follow these steps to configure your switch as a candidate BSR. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip pim bsr-candidate <i>interface-id</i> <i>hash-mask-length</i> [<i>priority</i>]	Configure your switch to be a candidate BSR. - For <i>interface-id</i>, enter the interface on this switch from which the BSR address is derived to make it a candidate. This interface must be enabled with PIM. Valid interfaces include physical ports, port channels, and VLANs. - For <i>hash-mask-length</i>, specify the mask length (32 bits maximum) that is to be ANDed with the group address before the hash function is called. All groups with the same seed hash correspond to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. (Optional) For <i>priority</i>, enter a number from 0 to 255. The BSR with the larger priority is preferred. If the priority values are the same, the device with the highest IP address is selected as the BSR. The default is 0.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove this device as a candidate BSR, use the **no ip pim bsr-candidate** global configuration command.

This example shows how to configure a candidate BSR, which uses the IP address 172.21.24.18 on a port as the advertised BSR address, uses 30 bits as the hash-mask-length, and has a priority of 10.

```
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# ip address 172.21.24.18 255.255.255.0
Switch(config-if)# ip pim sparse-dense-mode
Switch(config-if)# ip pim bsr-candidate gigabitethernet1/0/2 30 10
```

Configuring Candidate RPs

You can configure one or more candidate RPs. Similar to BSRs, the RPs should also have good connectivity to other devices and be in the backbone portion of the network. An RP can serve the entire IP multicast address space or an interface on it. Candidate RPs send candidate RP advertisements to the BSR. When deciding which devices should be RPs, consider these options:

- In a network of Cisco routers and multilayer switches where only Auto-RP is used, any device can be configured as an RP.
- In a network that includes only Cisco PIMv2 routers and multilayer switches and with routers from other vendors, any device can be used as an RP.
- In a network of Cisco PIMv1 routers, Cisco PIMv2 routers, and routers from other vendors, configure only Cisco PIMv2 routers and multilayer switches as RPs.

Beginning in privileged EXEC mode, follow these steps to configure your switch to advertise itself as a PIMv2 candidate RP to the BSR. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip pim rp-candidate <i>interface-id</i> [group-list <i>access-list-number</i>]	Configure your switch to be a candidate RP. • For <i>interface-id</i>, specify the interface whose associated IP address is advertised as a candidate RP address. Valid interfaces include physical ports, port channels, and VLANs. • (Optional) For group-list <i>access-list-number</i>, enter an IP standard access list number from 1 to 99. If no group-list is specified, the switch is a candidate RP for all groups.
Step 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. • 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 number of the network or host from which the packet is being sent. • (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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove this device as a candidate RP, use the **no ip pim rp-candidate** *interface-id* global configuration command.

This example shows how to configure the switch to advertise itself as a candidate RP to the BSR in its PIM domain. Standard access list number 4 specifies the group prefix associated with the RP that has the address identified by a port. That RP is responsible for the groups with the prefix 239.

```
Switch(config)# ip pim rp-candidate gigabitethernet1/0/2 group-list 4
Switch(config)# access-list 4 permit 239.0.0.0 0.255.255.255
```

Using Auto-RP and a BSR

If there are only Cisco devices in your network (no routers from other vendors), there is no need to configure a BSR. Configure Auto-RP in a network that is running both PIMv1 and PIMv2.

If you have non-Cisco PIMv2 routers that need to interoperate with Cisco PIMv1 routers and multilayer switches, both Auto-RP and a BSR are required. We recommend that a Cisco PIMv2 router or multilayer switch be both the Auto-RP mapping agent and the BSR.

If you must have one or more BSRs, we have these recommendations:

- Configure the candidate BSRs as the RP-mapping agents for Auto-RP. For more information, see the [“Configuring Auto-RP” section on page 31-13](#) and the [“Configuring Candidate BSRs” section on page 31-19](#).
- For group prefixes advertised through Auto-RP, the PIMv2 BSR mechanism should not advertise a subrange of these group prefixes served by a different set of RPs. In a mixed PIMv1 and PIMv2 domain, have backup RPs serve the same group prefixes. This prevents the PIMv2 DRs from selecting a different RP from those PIMv1 DRs, due to the longest match lookup in the RP-mapping database.

Beginning in privileged EXEC mode, follow these steps to verify the consistency of group-to-RP mappings. This procedure is optional.

	Command	Purpose
Step 1	show ip pim rp <i>[[group-name group-address] mapping]</i>	On any Cisco device, display the available RP mappings. • (Optional) For <i>group-name</i>, specify the name of the group about which to display RPs. • (Optional) For <i>group-address</i>, specify the address of the group about which to display RPs. • (Optional) Use the mapping keyword to display all group-to-RP mappings of which the Cisco device is aware (either configured or learned from Auto-RP).
Step 2	show ip pim rp-hash <i>group</i>	On a PIMv2 router or multilayer switch, confirm that the same RP is the one that a PIMv1 system chooses. For <i>group</i> , enter the group address for which to display RP information.

Monitoring the RP Mapping Information

To monitor the RP mapping information, use these commands in privileged EXEC mode:

- **show ip pim bsr** displays information about the elected BSR.
- **show ip pim rp-hash group** displays the RP that was selected for the specified group.
- **show ip pim rp [group-name | group-address | mapping]** displays how the switch learns of the RP (through the BSR or the Auto-RP mechanism).

Troubleshooting PIMv1 and PIMv2 Interoperability Problems

When debugging interoperability problems between PIMv1 and PIMv2, check these in the order shown:

1. Verify RP mapping with the **show ip pim rp-hash** privileged EXEC command, making sure that all systems agree on the same RP for the same group.
2. Verify interoperability between different versions of DRs and RPs. Make sure the RPs are interacting with the DRs properly (by responding with register-stops and forwarding decapsulated data packets from registers).


Configuring Advanced PIM Features

These sections describe the optional advanced PIM features:

- [Understanding PIM Shared Tree and Source Tree, page 31-22](#)
- [Delaying the Use of PIM Shortest-Path Tree, page 31-24](#) (optional)
- [Modifying the PIM Router-Query Message Interval, page 31-25](#) (optional)

Understanding PIM Shared Tree and Source Tree

By default, members of a group receive data from senders to the group across a single data-distribution tree rooted at the RP. [Figure 31-4](#) shows this type of shared-distribution tree. Data from senders is delivered to the RP for distribution to group members joined to the shared tree.

Figure 31-4 Shared Tree and Source Tree (Shortest-Path Tree)

If the data rate warrants, leaf routers (routers without any downstream connections) on the shared tree can use the data distribution tree rooted at the source. This type of distribution tree is called a shortest-path tree or source tree. By default, the software switches to a source tree upon receiving the first data packet from a source.

This process describes the move from a shared tree to a source tree:

1. A receiver joins a group; leaf Router C sends a join message toward the RP.
2. The RP puts a link to Router C in its outgoing interface list.
3. A source sends data; Router A encapsulates the data in a register message and sends it to the RP.
4. The RP forwards the data down the shared tree to Router C and sends a join message toward the source. At this point, data might arrive twice at Router C, once encapsulated and once natively.
5. When data arrives natively (unencapsulated) at the RP, it sends a register-stop message to Router A.
6. By default, reception of the first data packet prompts Router C to send a join message toward the source.
7. When Router C receives data on (S,G), it sends a prune message for the source up the shared tree.
8. The RP deletes the link to Router C from the outgoing interface of (S,G). The RP triggers a prune message toward the source.

Join and prune messages are sent for sources and RPs. They are sent hop-by-hop and are processed by each PIM device along the path to the source or RP. Register and register-stop messages are not sent hop-by-hop. They are sent by the designated router that is directly connected to a source and are received by the RP for the group.

Multiple sources sending to groups use the shared tree.

You can configure the PIM device to stay on the shared tree. For more information, see the [“Delaying the Use of PIM Shortest-Path Tree”](#) section on page 31-24.

Delaying the Use of PIM Shortest-Path Tree

The change from shared to source tree happens when the first data packet arrives at the last-hop router (Router C in [Figure 31-4](#)). This change occurs because the **ip pim spt-threshold** global configuration command controls that timing.

The shortest-path tree requires more memory than the shared tree but reduces delay. You might want to postpone its use. Instead of allowing the leaf router to immediately move to the shortest-path tree, you can specify that the traffic must first reach a threshold.

You can configure when a PIM leaf router should join the shortest-path tree for a specified group. If a source sends at a rate greater than or equal to the specified kbps rate, the multilayer switch triggers a PIM join message toward the source to construct a source tree (shortest-path tree). If the traffic rate from the source drops below the threshold value, the leaf router switches back to the shared tree and sends a prune message toward the source.

You can specify to which groups the shortest-path tree threshold applies by using a group list (a standard access list). If a value of 0 is specified or if the group list is not used, the threshold applies to all groups.

Beginning in privileged EXEC mode, follow these steps to configure a traffic rate threshold that must be reached before multicast routing is switched from the source tree to the shortest-path tree. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	access-list <i>access-list-number</i> { deny permit } <i>source</i> [<i>source-wildcard</i>]	Create a standard access list. - For <i>access-list-number</i>, the range is 1 to 99. - The deny keyword denies access if the conditions are matched. The permit keyword permits access if the conditions are matched. - For <i>source</i>, specify the multicast group to which the threshold will apply. (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.
Step 3	ip pim spt-threshold { <i>kbps</i> infinity } [group-list <i>access-list-number</i>]	Specify the threshold that must be reached before moving to shortest-path tree (spt). - For <i>kbps</i>, specify the traffic rate in kbps. The default is 0. Note Because of Catalyst 3750 Metro switch hardware limitations, 0 kbps is the only valid entry even though the range is 0 to 4294967. - Specify infinity if you want all sources for the specified group to use the shared tree, never switching to the source tree. (Optional) For group-list <i>access-list-number</i>, specify the access list created in Step 2. If the value is 0 or if the group-list is not used, the threshold applies to all groups.

	Command	Purpose
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no ip pim spt-threshold {kpbs | infinity}** global configuration command.

Modifying the PIM Router-Query Message Interval

PIM routers and multilayer switches send PIM router-query messages to find which device will be the DR for each LAN segment (subnet). The DR is responsible for sending IGMP host-query messages to all hosts on the directly connected LAN.

With PIM DM operation, the DR has meaning only if IGMPv1 is in use. IGMPv1 does not have an IGMP querier election process, so the elected DR functions as the IGMP querier. With PIM SM operation, the DR is the device that is directly connected to the multicast source. It sends PIM register messages to notify the RP that multicast traffic from a source needs to be forwarded down the shared tree. In this case, the DR is the device with the highest IP address.

Beginning in privileged EXEC mode, follow these steps to modify the router-query message interval. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to be configured, and enter interface configuration mode.
Step 3	ip pim query-interval <i>seconds</i>	Configure the frequency at which the switch sends PIM router-query messages. The default is 30 seconds. The range is 1 to 65535.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip igmp interface [<i>interface-id</i>]	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no ip pim query-interval [seconds]** interface configuration command.

Configuring Optional IGMP Features

These sections describe how to configure optional IGMP features:

- [Default IGMP Configuration, page 31-26](#)
- [Configuring the Switch as a Member of a Group, page 31-26](#) (optional)
- [Controlling Access to IP Multicast Groups, page 31-27](#) (optional)
- [Changing the IGMP Version, page 31-28](#) (optional)
- [Modifying the IGMP Host-Query Message Interval, page 31-29](#) (optional)
- [Changing the IGMP Query Timeout for IGMPv2, page 31-30](#) (optional)
- [Changing the Maximum Query Response Time for IGMPv2, page 31-30](#) (optional)
- [Configuring the Switch as a Statically Connected Member, page 31-31](#) (optional)

Default IGMP Configuration

[Table 31-3](#) shows the default IGMP configuration.

Table 31-3 Default IGMP Configuration

Feature	Default Setting
Multilayer switch as a member of a multicast group	No group memberships are defined.
Access to multicast groups	All groups are allowed on an interface.
IGMP version	Version 2 on all interfaces.
IGMP host-query message interval	60 seconds on all interfaces.
IGMP query timeout	60 seconds on all interfaces.
IGMP maximum query response time	10 seconds on all interfaces.
Multilayer switch as a statically connected member	Disabled.

Configuring the Switch as a Member of a Group

You can configure the switch as a member of a multicast group and discover multicast reachability in a network. If all the multicast-capable routers and multilayer switches that you administer are members of a multicast group, pinging that group causes all these devices to respond. The devices respond to ICMP echo-request packets addressed to a group of which they are members. Another example is the multicast trace-route tools provided in the software.



Caution

Performing this procedure might impact the CPU performance because the CPU will receive all data traffic for the group address.

Beginning in privileged EXEC mode, follow these steps to configure the switch to be a member of a group. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to be configured, and enter interface configuration mode.
Step 3	ip igmp join-group <i>group-address</i>	Configure the switch to join a multicast group. By default, no group memberships are defined. For <i>group-address</i> , specify the multicast IP address in dotted decimal notation.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip igmp interface [<i>interface-id</i>]	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To cancel membership in a group, use the **no ip igmp join-group** *group-address* interface configuration command.

This example shows how to enable the switch to join multicast group 255.2.2.2:

```
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# ip igmp join-group 255.2.2.2
```

Controlling Access to IP Multicast Groups

The switch sends IGMP host-query messages to find which multicast groups have members on attached local networks. The switch then forwards to these group members all packets addressed to the multicast group. You can place a filter on each interface to restrict the multicast groups that hosts on the subnet serviced by the interface can join.

Beginning in privileged EXEC mode, follow these steps to filter multicast groups allowed on an interface. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to be configured, and enter interface configuration mode.
Step 3	ip igmp access-group <i>access-list-number</i>	Specify the multicast groups that hosts on the subnet serviced by an interface can join. By default, all groups are allowed on an interface. For <i>access-list-number</i> , specify an IP standard access list number. The range is 1 to 99.
Step 4	exit	Return to global configuration mode.

	Command	Purpose
Step 5	access-list <i>access-list-number</i> { deny permit } <i>source</i> [<i>source-wildcard</i>]	Create a standard access list. - For <i>access-list-number</i>, specify the access list created in Step 3. - The deny keyword denies access if the conditions are matched. The permit keyword permits access if the conditions are matched. - For <i>source</i>, specify the multicast group that hosts on the subnet can join. (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.
Step 6	end	Return to privileged EXEC mode.
Step 7	show ip igmp interface [<i>interface-id</i>]	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable groups on an interface, use the **no ip igmp access-group** interface configuration command.

This example shows how to configure hosts attached to a port as able to join only group 255.2.2.2:

```
Switch(config)# access-list 1 255.2.2.2 0.0.0.0
Switch(config-if)# interface gigabitethernet1/0/1
Switch(config-if)# ip igmp access-group 1
```

Changing the IGMP Version

By default, the switch uses IGMP Version 2, which provides features such as the IGMP query timeout and the maximum query response time.

All systems on the subnet must support the same version. The switch does not automatically detect Version 1 systems and switch to Version 1. You can mix Version 1 and Version 2 hosts on the subnet because Version 2 routers or switches always work correctly with IGMPv1 hosts.

Configure the switch for Version 1 if your hosts do not support Version 2.

Beginning in privileged EXEC mode, follow these steps to change the IGMP version. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to be configured, and enter interface configuration mode.
Step 3	ip igmp version {1 2}	Specify the IGMP version that the switch uses. Note If you change to Version 1, you cannot configure the ip igmp query-interval or the ip igmp query-max-response-time interface configuration commands.
Step 4	end	Return to privileged EXEC mode.

	Command	Purpose
Step 5	show ip igmp interface [<i>interface-id</i>]	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no ip igmp version** interface configuration command.

Modifying the IGMP Host-Query Message Interval

The switch periodically sends IGMP host-query messages to discover which multicast groups are present on attached networks. These messages are sent to the all-hosts multicast group (224.0.0.1) with a time-to-live (TTL) of 1. The switch sends host-query messages to refresh its knowledge of memberships present on the network. If, after some number of queries, the software discovers that no local hosts are members of a multicast group, the software stops forwarding multicast packets to the local network from remote origins for that group and sends a prune message upstream toward the source.

The switch elects a PIM designated router (DR) for the LAN (subnet). The DR is the router or multilayer switch with the highest IP address for IGMPv2. For IGMPv1, the DR is elected according to the multicast routing protocol that runs on the LAN. The designated router is responsible for sending IGMP host-query messages to all hosts on the LAN. In sparse mode, the designated router also sends PIM register and PIM join messages toward the RP router.

Beginning in privileged EXEC mode, follow these steps to modify the host-query interval. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to be configured, and enter interface configuration mode.
Step 3	ip igmp query-interval <i>seconds</i>	Configure the frequency at which the designated router sends IGMP host-query messages. By default, the designated router sends IGMP host-query messages every 60 seconds to keep the IGMP overhead very low on hosts and networks. The range is 1 to 65535.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip igmp interface [<i>interface-id</i>]	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no ip igmp query-interval** interface configuration command.

Changing the IGMP Query Timeout for IGMPv2

If you are using IGMPv2, you can specify the period of time before the switch takes over as the querier for the interface. By default, the switch waits twice the query interval controlled by the **ip igmp query-interval** interface configuration command. After that time, if the switch has received no queries, it becomes the querier.

You can configure the query interval by entering the **show ip igmp interface interface-id** privileged EXEC command.

Beginning in privileged EXEC mode, follow these steps to change the IGMP query timeout. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface interface-id	Specify the interface to be configured, and enter interface configuration mode.
Step 3	ip igmp querier-timeout seconds	Specify the IGMP query timeout. The default is 60 seconds (twice the query interval). The range is 60 to 300.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip igmp interface [interface-id]	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no ip igmp querier-timeout** interface configuration command.

Changing the Maximum Query Response Time for IGMPv2

If you are using IGMPv2, you can change the maximum query response time advertised in IGMP queries. The maximum query response time enables the switch to quickly detect that there are no more directly connected group members on a LAN. Decreasing the value enables the switch to prune groups faster.

Beginning in privileged EXEC mode, follow these steps to change the maximum query response time. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface interface-id	Specify the port to be configured, and enter interface configuration mode.
Step 3	ip igmp query-max-response-time seconds	Change the maximum query response time advertised in IGMP queries. The default is 10 seconds. The range is 1 to 25.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip igmp interface [interface-id]	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no ip igmp query-max-response-time** interface configuration command.

Configuring the Switch as a Statically Connected Member

Sometimes there is either no group member on a network segment or a host cannot report its group membership by using IGMP. However, you might want multicast traffic to go to that network segment. These are ways to pull multicast traffic down to a network segment:

- Use the **ip igmp join-group** interface configuration command. With this method, the switch accepts the multicast packets in addition to forwarding them. Accepting the multicast packets prevents the switch from fast switching.
- Use the **ip igmp static-group** interface configuration command. With this method, the switch does not accept the packets itself, but only forwards them. This method enables fast switching. The outgoing interface appears in the IGMP cache, but the switch itself is not a member, as evidenced by lack of an *L* (local) flag in the multicast route entry.

Beginning in privileged EXEC mode, follow these steps to configure the switch itself to be a statically connected member of a group (and enable fast switching). This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to be configured, and enter interface configuration mode.
Step 3	ip igmp static-group <i>group-address</i>	Configure the switch as a statically connected member of a group. By default, this feature is disabled.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip igmp interface [<i>interface-id</i>]	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the switch as a member of the group, use the **no ip igmp static-group** *group-address* interface configuration command.

Configuring Optional Multicast Routing Features

This section describes how to configure optional multicast routing features, which are grouped as follows:

- Features for Layer 2 connectivity and MBONE multimedia conference session and set up:
 - [Enabling CGMP Server Support, page 31-32](#) (optional)
 - [Configuring sdr Listener Support, page 31-33](#) (optional)
- Features that control bandwidth utilization:
 - [Configuring an IP Multicast Boundary, page 31-34](#) (optional)

Enabling CGMP Server Support

The switch serves as a CGMP server for devices that do not support IGMP snooping but have CGMP client functionality. CGMP is a protocol used on Cisco routers and multilayer switches connected to Layer 2 Catalyst switches to perform tasks similar to those performed by IGMP. CGMP is necessary because the Layer 2 switch cannot distinguish between IP multicast data packets and IGMP report messages, which are both at the MAC-level and are addressed to the same group address.

Beginning in privileged EXEC mode, follow these steps to enable the CGMP server on the switch interface. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface that is connected to the Layer 2 Catalyst switch, and enter interface configuration mode.
Step 3	ip cgmp [proxy]	Enable CGMP on the interface. By default, CGMP is disabled on all interfaces. Enabling CGMP triggers a CGMP join message. Enable CGMP only on Layer 3 interfaces connected to Layer 2 Catalyst switches. (Optional) When you enter the proxy keyword, the CGMP proxy function is enabled. The proxy router advertises the existence of non-CGMP-capable routers by sending a CGMP join message with the non-CGMP-capable router MAC address and a group address of 0000.0000.0000. Note To perform CGMP proxy, the switch must be the IGMP querier. If you configure the ip cgmp proxy command, you must manipulate the IP addresses so that the switch is the IGMP querier, which might be the highest or lowest IP address, depending on which version of IGMP is running on the network. An IGMP Version 2 querier is selected based on the lowest IP address on the interface. An IGMP Version 1 querier is selected based on the multicast routing protocol used on the interface.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.
Step 7		Verify the Layer 2 Catalyst switch CGMP-client configuration. For more information, refer to the documentation that shipped with the product.

To disable CGMP on the interface, use the **no ip cgmp** interface configuration command.

When multiple Cisco CGMP-capable devices are connected to a switched network and the **ip cgmp proxy** command is needed, we recommend that all devices be configured with the same CGMP option and have precedence for becoming the IGMP querier over non-Cisco routers.

Configuring sdr Listener Support

The MBONE is the small subset of Internet routers and hosts that are interconnected and capable of forwarding IP multicast traffic. Other interesting multimedia content is often broadcast over the MBONE. Before you can join a multimedia session, you need to know what multicast group address and interface are being used for the session, when the session is going to be active, and what sort of applications (audio, video, and so forth) are required on your workstation. The MBONE Session Directory version 2 (sdr) tool provides this information. This freeware application can be downloaded from several sites on the World Wide Web, one of which is <http://www.video.ja.net/mice/index.html>.

SDR is a multicast application that listens to a well-known multicast group address and interface for Session Announcement Protocol (SAP) multicast packets from SAP clients, which announce their conference sessions. These SAP packets contain a session description, the time the session is active, its IP multicast group addresses, media format, contact person, and other information about the advertised multimedia session. The information in the SAP packet appears in the SDR Session Announcement window.

Enabling sdr Listener Support

By default, the switch does not listen to session directory advertisements.

Beginning in privileged EXEC mode, follow these steps to enable the switch to join the default session directory group (224.2.127.254) on the interface and listen to session directory advertisements. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to be enabled for sdr, and enter interface configuration mode.
Step 3	ip sdr listen	Enable sdr listener supinterface.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable sdr support, use the **no ip sdr listen** interface configuration command.

Limiting How Long an sdr Cache Entry Exists

By default, entries are never deleted from the sdr cache. You can limit how long the entry remains active so that if a source stops advertising SAP information, old advertisements are not needlessly kept.

Beginning in privileged EXEC mode, follow these steps to limit how long an sdr cache entry stays active in the cache. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip sdr cache-timeout <i>minutes</i>	Limit how long an sdr cache entry stays active in the cache. By default, entries are never deleted from the cache. For <i>minutes</i> , the range is 1 to 4294967295.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no ip sdr cache-timeout** global configuration command. To delete the entire cache, use the **clear ip sdr** privileged EXEC command.

To display the session directory cache, use the **show ip sdr** privileged EXEC command.

Configuring an IP Multicast Boundary


Administratively-scoped boundaries can be used to limit the forwarding of multicast traffic outside of a domain or subdomain. This approach uses a special range of multicast addresses, called *administratively-scoped addresses*, as the boundary mechanism. If you configure an administratively-scoped boundary on a routed port, multicast traffic whose multicast group addresses fall in this range can not enter or exit this port, thereby providing a firewall for multicast traffic in this address range.



Note

Multicast boundaries and TTL thresholds control the scoping of multicast domains; however, TTL thresholds are not supported by the switch. You should use multicast boundaries instead of TTL thresholds to limit the forwarding of multicast traffic outside of a domain or a subdomain.

Figure 31-5 shows that Company XYZ has an administratively-scoped boundary set for the multicast address range 239.0.0.0/8 on all routed ports at the perimeter of its network. This boundary prevents any multicast traffic in the range 239.0.0.0 through 239.255.255.255 from entering or leaving the network. Similarly, the engineering and marketing departments have an administratively-scoped boundary of 239.128.0.0/16 around the perimeter of their networks. This boundary prevents multicast traffic in the range of 239.128.0.0 through 239.128.255.255 from entering or leaving their respective networks.

Figure 31-5 Administratively-Scoped Boundaries

You can define an administratively-scoped boundary on a routed port for multicast group addresses. A standard access list defines the range of addresses affected. When a boundary is defined, no multicast data packets are allowed to flow across the boundary from either direction. The boundary allows the same multicast group address to be reused in different administrative domains.

The IANA has designated the multicast address range 239.0.0.0 to 239.255.255.255 as the administratively-scoped addresses. This range of addresses can then be reused in domains administered by different organizations. The addresses would be considered local, not globally unique.

Beginning in privileged EXEC mode, follow these steps to set up an administratively-scoped boundary. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	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. - For <i>access-list-number</i>, the range is 1 to 99. - 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 number of the network or host from which the packet is being sent. (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.
Step 3	interface <i>interface-id</i>	Specify the interface to be configured, and enter interface configuration mode.
Step 4	ip multicast boundary <i>access-list-number</i>	Configure the boundary, specifying the access list you created in Step 2.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the boundary, use the **no ip multicast boundary** interface configuration command.

This example shows how to set up a boundary for all administratively-scoped addresses:

```
Switch(config)# access-list 1 deny 239.0.0.0 0.255.255.255
Switch(config)# access-list 1 permit 224.0.0.0 15.255.255.255
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# ip multicast boundary 1
```

Configuring Basic DVMRP Interoperability Features

These sections describe how to perform basic configuration tasks on your switch to interoperate with DVMRP devices:

- [Configuring DVMRP Interoperability, page 31-36](#) (optional)
- [Configuring a DVMRP Tunnel, page 31-38](#) (optional)
- [Advertising Network 0.0.0.0 to DVMRP Neighbors, page 31-40](#) (optional)
- [Responding to mroute Requests, page 31-41](#) (optional)

For more advanced DVMRP features, see the [“Configuring Advanced DVMRP Interoperability Features”](#) section on page 31-41.

Configuring DVMRP Interoperability

Cisco multicast routers and multilayer switches using PIM can interoperate with non-Cisco multicast routers that use the DVMRP.

PIM devices dynamically discover DVMRP multicast routers on attached networks by listening to DVMRP probe messages. When a DVMRP neighbor has been discovered, the PIM device periodically sends DVMRP report messages advertising the unicast sources reachable in the PIM domain. By default, directly connected subnets and networks are advertised. The device forwards multicast packets that have been forwarded by DVMRP routers and, in turn, forwards multicast packets to DVMRP routers.

You can configure an access list on the PIM routed port connected to the MBONE to limit the number of unicast routes that are advertised in DVMRP route reports. Otherwise, all routes in the unicast routing table are advertised.



Note

The mroute protocol is a public-domain implementation of DVMRP. You must use mroute Version 3.8 (which implements a nonpruning version of DVMRP) when Cisco routers and multilayer switches are directly connected to DVMRP routers or interoperate with DVMRP routers over an MBONE tunnel. DVMRP advertisements produced by the Cisco IOS software can cause older versions of the mroute protocol to corrupt their routing tables and those of their neighbors.

You can configure what sources are advertised and what metrics are used by configuring the **ip dvmrp metric** interface configuration command. You can also direct all sources learned through a particular unicast routing process to be advertised into DVMRP.

Beginning in privileged EXEC mode, follow these steps to configure the sources that are advertised and the metrics that are used when DVMRP route-report messages are sent. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	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. - For <i>access-list-number</i>, the range is 1 to 99. - 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 number of the network or host from which the packet is being sent. (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.
Step 3	interface <i>interface-id</i>	Specify the interface connected to the MBONE and enabled for multicast routing, and enter interface configuration mode.
Step 4	ip dvmrp metric <i>metric</i> [list <i>access-list-number</i>] [[<i>protocol process-id</i>] [dvmrp]]	Configure the metric associated with a set of destinations for DVMRP reports. - For <i>metric</i>, the range is 0 to 32. A value of 0 means that the route is not advertised. A value of 32 is equivalent to infinity (unreachable). (Optional) For list <i>access-list-number</i>, enter the access list number created in Step 2. If specified, only the multicast destinations that match the access list are reported with the configured metric. (Optional) For <i>protocol process-id</i>, enter the name of the unicast routing protocol, such as eigrp, igrp, ospf, rip, static, or dvmrp, and the process ID number of the routing protocol. If specified, only routes learned by the specified routing protocol are advertised in DVMRP report messages. (Optional) If specified, the dvmrp keyword allows routes from the DVMRP routing table to be advertised with the configured <i>metric</i> or filtered.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the metric or route map, use the **no ip dvmrp metric** *metric* [**list** *access-list-number*] [[*protocol process-id*] | [**dvmrp**]] or the **no ip dvmrp metric** *metric* **route-map** *map-name* interface configuration command.

A more sophisticated way to achieve the same results as the preceding command is to use a route map (**ip dvmrp metric** *metric* **route-map** *map-name* interface configuration command) instead of an access list. You subject unicast routes to route-map conditions before they are injected into DVMRP.

This example shows how to configure DVMRP interoperability when the PIM device and the DVMRP router are on the same network segment. In this example, access list 1 advertises the networks (198.92.35.0, 198.92.36.0, 198.92.37.0, 131.108.0.0, and 150.136.0.0) to the DVMRP router, and access list 2 prevents all other networks from being advertised (**ip dvmrp metric 0** interface configuration command).

```
Switch(config-if)# interface gigabitethernet1/0/1
Switch(config-if)# ip address 131.119.244.244 255.255.255.0
Switch(config-if)# ip pim dense-mode
Switch(config-if)# ip dvmrp metric 1 list 1
Switch(config-if)# ip dvmrp metric 0 list 2
Switch(config-if)# exit
Switch(config)# access-list 1 permit 198.92.35.0 0.0.0.255
Switch(config)# access-list 1 permit 198.92.36.0 0.0.0.255
Switch(config)# access-list 1 permit 198.92.37.0 0.0.0.255
Switch(config)# access-list 1 permit 131.108.0.0 0.0.255.255
Switch(config)# access-list 1 permit 150.136.0.0 0.0.255.255
Switch(config)# access-list 1 deny 0.0.0.0 255.255.255.255
Switch(config)# access-list 2 permit 0.0.0.0 255.255.255.255
```

Configuring a DVMRP Tunnel

The software supports DVMRP tunnels to the MBONE. You can configure a DVMRP tunnel on a router or multilayer switch if the other end is running DVMRP. The software then sends and receives multicast packets through the tunnel. This strategy enables a PIM domain to connect to the DVMRP router when all routers on the path do not support multicast routing. You cannot configure a DVMRP tunnel between two routers.

When a Cisco router or multilayer switch runs DVMRP through a tunnel, it advertises sources in DVMRP report messages, much as it does on real networks. The software also caches DVMRP report messages it receives and uses them in its RPF calculation. This behavior enables the software to forward multicast packets received through the tunnel.

When you configure a DVMRP tunnel, you should assign an IP address to a tunnel in these cases:

- To send IP packets through the tunnel
- To configure the software to perform DVMRP summarization

The software does not advertise subnets through the tunnel if the tunnel has a different network number from the subnet. In this case, the software advertises only the network number through the tunnel.

Beginning in privileged EXEC mode, follow these steps to configure a DVMRP tunnel. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	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. - For <i>access-list-number</i>, the range is 1 to 99. - 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 number of the network or host from which the packet is being sent. (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.
Step 3	interface tunnel <i>number</i>	Specify a tunnel port, and enter interface configuration mode.
Step 4	tunnel source <i>ip-address</i>	Specify the source address of the tunnel port. Enter the IP address of the interface on the switch.
Step 5	tunnel destination <i>ip-address</i>	Specify the destination address of the tunnel port. Enter the IP address of the mrouter.
Step 6	tunnel mode dvmrp	Configure the encapsulation mode for the tunnel to DVMRP.
Step 7	ip address <i>address mask</i> or ip unnumbered <i>type number</i>	Assign an IP address to the interface. or Configure the interface as unnumbered.
Step 8	ip pim [dense-mode sparse-mode]	Configure the PIM mode on the interface.
Step 9	ip dvmrp accept-filter <i>access-list-number</i> [<i>distance</i>] neighbor-list <i>access-list-number</i>	Configure an acceptance filter for incoming DVMRP reports. By default, all destination reports are accepted with a distance of 0. Reports from all neighbors are accepted. - For <i>access-list-number</i>, specify the access list number created in Step 2. Any sources that match the access list are stored in the DVMRP routing table with distance. (Optional) For <i>distance</i>, enter the administrative distance to the destination. By default, the administrative distance for DVMRP routes is 0 and take precedence over unicast routing table routes. If you have two paths to a source, one through unicast routing (using PIM as the multicast routing protocol) and another using DVMRP, and if you want to use the PIM path, increase the administrative distance for DVMRP routes. The range is 1 to 255. - For neighbor-list <i>access-list-number</i>, enter the number of the neighbor list created in Step 2. DVMRP reports are accepted only by those neighbors on the list.
Step 10	end	Return to privileged EXEC mode.

	Command	Purpose
Step 11	show running-config	Verify your entries.
Step 12	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the filter, use the **no ip dvmrp accept-filter** *access-list-number* [*distance*] **neighbor-list** *access-list-number* interface configuration command.

This example shows how to configure a DVMRP tunnel. In this configuration, the IP address of the tunnel on the Cisco switch is assigned *unnumbered*, which causes the tunnel to appear to have the same IP address as port 1. The tunnel endpoint source address is 172.16.2.1, and the tunnel endpoint address of the remote DVMRP router to which the tunnel is connected is 192.168.1.10. Any packets sent through the tunnel are encapsulated in an outer IP header. The Cisco switch is configured to accept incoming DVMRP reports with a distance of 100 from 198.92.37.0 through 198.92.37.255.

```
Switch(config)# ip multicast-routing
Switch(config)# interface tunnel 0
Switch(config-if)# ip unnumbered gigabitethernet1/0/1
Switch(config-if)# ip pim dense-mode
Switch(config-if)# tunnel source gigabitethernet1/0/1
Switch(config-if)# tunnel destination 192.168.1.10
Switch(config-if)# tunnel mode dvmrp
Switch(config-if)# ip dvmrp accept-filter 1 100
Switch(config-if)# interface gigabitethernet1/0/1
Switch(config-if)# ip address 172.16.2.1 255.255.255.0
Switch(config-if)# ip pim dense-mode
Switch(config)# exit
Switch(config)# access-list 1 permit 198.92.37.0 0.0.0.255
```

Advertising Network 0.0.0.0 to DVMRP Neighbors

If your switch is a neighbor of an mrouted version 3.6 device, you can configure the software to advertise network 0.0.0.0 (the default route) to the DVMRP neighbor. The DVMRP default route computes the RPF information for any multicast sources that do not match a more specific route.

Do not advertise the DVMRP default into the MBONE.

Beginning in privileged EXEC mode, follow these steps to advertise network 0.0.0.0 to DVMRP neighbors on an interface. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface that is connected to the DVMRP router, and enter interface configuration mode.
Step 3	ip dvmrp default-information {originate only}	Advertise network 0.0.0.0 to DVMRP neighbors. Use this command only when the switch is a neighbor of mrouted version 3.6 machines. The keywords have these meanings: • originate—Specifies that other routes more specific than 0.0.0.0 can also be advertised. • only—Specifies that no DVMRP routes other than 0.0.0.0 are advertised.

	Command	Purpose
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To prevent the default route advertisement, use the **no ip dvmrp default-information** interface configuration command.

Responding to mrinfo Requests

The software answers mrinfo requests sent by mrouted systems and Cisco routers and multilayer switches. The software returns information about neighbors through DVMRP tunnels and all the routed ports. This information includes the metric (always set to 1), the configured TTL threshold, the status of the port, and various flags. You can also use the **mrinfo** privileged EXEC command to query the router or switch itself, as in this example:

```
Switch# mrinfo
171.69.214.27 (mm1-7kd.cisco.com) [version cisco 11.1] [flags: PMS]:
171.69.214.27 -> 171.69.214.26 (mm1-r7kb.cisco.com) [1/0/pim/querier]
171.69.214.27 -> 171.69.214.25 (mm1-45a.cisco.com) [1/0/pim/querier]
171.69.214.33 -> 171.69.214.34 (mm1-45c.cisco.com) [1/0/pim]
171.69.214.137 -> 0.0.0.0 [1/0/pim/querier/down/leaf]
171.69.214.203 -> 0.0.0.0 [1/0/pim/querier/down/leaf]
171.69.214.18 -> 171.69.214.20 (mm1-45e.cisco.com) [1/0/pim]
171.69.214.18 -> 171.69.214.19 (mm1-45c.cisco.com) [1/0/pim]
171.69.214.18 -> 171.69.214.17 (mm1-45a.cisco.com) [1/0/pim]
```

Configuring Advanced DVMRP Interoperability Features

Cisco routers and multilayer switches run PIM to forward multicast packets to receivers and receive multicast packets from senders. It is also possible to propagate DVMRP routes into and through a PIM cloud. PIM uses this information; however, Cisco routers and multilayer switches do not implement DVMRP to forward multicast packets.

These sections describe how to perform advanced optional configuration tasks on your switch to interoperate with DVMRP devices:

- [Enabling DVMRP Unicast Routing, page 31-42](#) (optional)
- [Rejecting a DVMRP Nonpruning Neighbor, page 31-43](#) (optional)
- [Controlling Route Exchanges, page 31-45](#) (optional)

For information on basic DVMRP features, see the “[Configuring Basic DVMRP Interoperability Features](#)” section on page 31-36.

Enabling DVMRP Unicast Routing

Because multicast routing and unicast routing require separate topologies, PIM must follow the multicast topology to build loopless distribution trees. Using DVMRP unicast routing, Cisco routers, multilayer switches, and mrouted-based machines exchange DVMRP unicast routes, to which PIM can then reverse-path forward.

Cisco devices do not perform DVMRP multicast routing among each other, but they can exchange DVMRP routes. The DVMRP routes provide a multicast topology that might differ from the unicast topology. This enables PIM to run over the multicast topology, thereby enabling sparse-mode PIM over the MBONE topology.

When DVMRP unicast routing is enabled, the router or switch caches routes learned in DVMRP report messages in a DVMRP routing table. When PIM is running, these routes might be preferred over routes in the unicast routing table, enabling PIM to run on the MBONE topology when it is different from the unicast topology.

DVMRP unicast routing can run on all interfaces. For DVMRP tunnels, it uses DVMRP multicast routing. This feature does not enable DVMRP multicast routing among Cisco routers and multilayer switches. However, if there is a DVMRP-capable multicast router, the Cisco device can do PIM/DVMRP multicast routing.

Beginning in privileged EXEC mode, follow these steps to enable DVMRP unicast routing. This procedure is optional.


	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface that is connected to the DVMRP router, and enter interface configuration mode.
Step 3	ip dvmrp unicast-routing	Enable DVMRP unicast routing (to send and receive DVMRP routes). This feature is disabled by default.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable this feature, use the **no ip dvmrp unicast-routing** interface configuration command.

Rejecting a DVMRP Nonpruning Neighbor


By default, Cisco devices accept all DVMRP neighbors as peers, regardless of their DVMRP capability. However, some non-Cisco devices run old versions of DVMRP that cannot prune, so they continuously receive forwarded packets, wasting bandwidth. [Figure 31-6](#) shows this scenario.

Figure 31-6 Leaf Nonpruning DVMRP Neighbor



101244

You can prevent the switch from peering (communicating) with a DVMRP neighbor if that neighbor does not support DVMRP pruning or grafting. To do so, configure the switch (which is a neighbor to the leaf, nonpruning DVMRP machine) with the **ip dvmrp reject-non-pruners** interface configuration command on the interface connected to the nonpruning machine as shown in [Figure 31-7](#). In this case, when the switch receives DVMRP probe or report message without the prune-capable flag set, the switch logs a syslog message and discards the message.

Figure 31-7 Router Rejects Nonpruning DVMRP Neighbor

Note that the **ip dvmrp reject-non-pruners** interface configuration command prevents peering with neighbors only. If there are any nonpruning routers multiple hops away (downstream toward potential receivers) that are not rejected, a nonpruning DVMRP network might still exist.

Beginning in privileged EXEC mode, follow these steps to prevent peering with nonpruning DVMRP neighbors. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface connected to the nonpruning DVMRP neighbor, and enter interface configuration mode.
Step 3	ip dvmrp reject-non-pruners	Prevent peering with nonpruning DVMRP neighbors.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable this function, use the **no ip dvmrp reject-non-pruners** interface configuration command.

Controlling Route Exchanges

These sections describe how to tune the Cisco device advertisements of DVMRP routes:

- [Limiting the Number of DVMRP Routes Advertised, page 31-45](#) (optional)
- [Changing the DVMRP Route Threshold, page 31-45](#) (optional)
- [Configuring a DVMRP Summary Address, page 31-46](#) (optional)
- [Disabling DVMRP Autosummarization, page 31-48](#) (optional)
- [Adding a Metric Offset to the DVMRP Route, page 31-48](#) (optional)

Limiting the Number of DVMRP Routes Advertised

By default, only 7000 DVMRP routes are advertised over an interface enabled to run DVMRP (that is, a DVMRP tunnel, an interface where a DVMRP neighbor has been discovered, or an interface configured to run the **ip dvmrp unicast-routing** interface configuration command).

Beginning in privileged EXEC mode, follow these steps to change the DVMRP route limit. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip dvmrp route-limit <i>count</i>	Change the number of DVMRP routes advertised over an interface enabled for DVMRP. This command prevents misconfigured ip dvmrp metric interface configuration commands from causing massive route injection into the MBONE. By default, 7000 routes are advertised. The range is 0 to 4294967295.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To configure no route limit, use the **no ip dvmrp route-limit** global configuration command.

Changing the DVMRP Route Threshold

By default, 10,000 DVMRP routes can be received per interface within a 1-minute interval. When that rate is exceeded, a syslog message is issued, warning that there might be a route surge occurring. The warning is typically used to quickly detect when devices have been misconfigured to inject a large number of routes into the MBONE.

Beginning in privileged EXEC mode, follow these steps to change the threshold number of routes that trigger the warning. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip dvmrp routehog-notification <i>route-count</i>	Configure the number of routes that trigger a syslog message. The default is 10,000 routes. The range is 1 to 4294967295.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting use the **no ip dvmrp routehog-notification** global configuration command.


Use the **show ip igmp interface** privileged EXEC command to display a running count of routes. When the count is exceeded, ***** ALERT ***** is appended to the line.

Configuring a DVMRP Summary Address

By default, a Cisco device advertises in DVMRP route-report messages only connected unicast routes (that is, only routes to subnets that are directly connected to the router) from its unicast routing table. These routes undergo normal DVMRP classful route summarization. This process depends on whether the route being advertised is in the same classful network as the interface over which it is being advertised.

Figure 31-8 shows an example of the default behavior. This example shows that the DVMRP report sent by the Cisco router contains the three original routes received from the DVMRP router that have been poison-reversed by adding 32 to the DVMRP metric. Listed after these routes are two routes that are advertisements for the two directly connected networks (176.32.10.0/24 and 176.32.15.0/24) that were taken from the unicast routing table. Because the DVMRP tunnel shares the same IP address as Fast Ethernet port 1 and falls into the same Class B network as the two directly connected subnets, classful summarization of these routes was not performed. As a result, the DVMRP router is able to poison-reverse only these two routes to the directly connected subnets and is able to only RPF properly for multicast traffic sent by sources on these two Ethernet segments. Any other multicast source in the network behind the Cisco router that is not on these two Ethernet segments does not properly RPF-check on the DVMRP router and is discarded.

You can force the Cisco router to advertise the summary address (specified by the address and mask pair in the **ip dvmrp summary-address address mask** interface configuration command) in place of any route that falls in this address range. The summary address is sent in a DVMRP route report if the unicast routing table contains at least one route in this range; otherwise, the summary address is not advertised. In Figure 31-8, you configure the **ip dvmrp summary-address** command on the Cisco router tunnel port. As a result, the Cisco router sends only a single summarized Class B advertisement for network 176.32.0.0.16 from the unicast routing table.

Figure 31-8 Only Connected Unicast Routes Are Advertised by Default

Beginning in privileged EXEC mode, follow these steps to customize the summarization of DVMRP routes if the default classful autosummarization does not suit your needs. This procedure is optional.

**Note**

At least one more-specific route must be present in the unicast routing table before a configured summary address is advertised.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface that is connected to the DVMRP router, and enter interface configuration mode.
Step 3	ip dvmrp summary-address <i>address mask</i> [metric value]	Specify a DVMRP summary address. - For summary-address <i>address mask</i>, specify the summary IP address and mask that is advertised instead of the more specific route. (Optional) For metric value, specify the metric that is advertised with the summary address. The default is 1. The range is 1 to 32.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the summary address, use the **no ip dvmrp summary-address** *address mask* [**metric value**] interface configuration command.

Disabling DVMRP Autosummarization

By default, the software automatically performs some level of DVMRP summarization. Disable this function if you want to advertise all routes, not just a summary. In some special cases, you can use the neighboring DVMRP router with all subnet information to better control the flow of multicast traffic in the DVMRP network. One such case might occur if the PIM network is connected to the DVMRP cloud at several points and more specific (unsummarized) routes are being injected into the DVMRP network to advertise better paths to individual subnets inside the PIM cloud.

If you configure the **ip dvmrp summary-address** interface configuration command and did not configure **no ip dvmrp auto-summary**, you get both custom and autosummaries.

Beginning in privileged EXEC mode, follow these steps to disable DVMRP autosummarization. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface connected to the DVMRP router, and enter interface configuration mode.
Step 3	no ip dvmrp auto-summary	Disable DVMRP autosummarization.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To re-enable auto summarization, use the **ip dvmrp auto-summary** interface configuration command.

Adding a Metric Offset to the DVMRP Route

By default, the switch increments by one the metric (hop count) of a DVMRP route advertised in incoming DVMRP reports. You can change the metric if you want to favor or not favor a certain route.

For example, a route is learned by multilayer switch A, and the same route is learned by multilayer switch B with a higher metric. If you want to use the path through switch B because it is a faster path, you can apply a metric offset to the route learned by switch A to make it larger than the metric learned by switch B, and you can choose the path through switch B.

Beginning in privileged EXEC mode, follow these steps to change the default metric. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to be configured, and enter interface configuration mode.

	Command	Purpose
Step 3	ip dvmrp metric-offset [in out] <i>increment</i>	Change the metric added to DVMRP routes advertised in incoming interfaces. The keywords have these meanings: (Optional) in—Specifies that the increment value is added to incoming DVMRP reports and is reported in mrinfo replies. (Optional) out—Specifies that the increment value is added to outgoing DVMRP reports for routes from the DVMRP routing table. If neither in nor out is specified, in is the default. For <i>increment</i>, specify the value that is added to the metric of a DVMRP router advertised in a report message. The range is 1 to 31. If the ip dvmrp metric-offset command is not configured on an interface, the default increment value for incoming routes is 1, and the default for outgoing routes is 0.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no ip dvmrp metric-offset** interface configuration command.

Monitoring and Maintaining IP Multicast Routing

These sections describe how to monitor and maintain IP multicast routing:

- [Clearing Caches, Tables, and Databases, page 31-49](#)
- [Displaying System and Network Statistics, page 31-50](#)
- [Monitoring IP Multicast Routing, page 31-51](#)

Clearing Caches, Tables, and Databases

You can remove all contents of a particular cache, table, or database. Clearing a cache, table, or database might be necessary when the contents of the particular structure are or suspected to be invalid.

You can use any of the privileged EXEC commands in [Table 31-4](#) to clear IP multicast caches, tables, and databases:

Table 31-4 Commands for Clearing Caches, Tables, and Databases

Command	Purpose
clear ip cgmp	Clear all group entries the Catalyst switches have cached.
clear ip dvmrp route { * <i>route</i> }	Delete routes from the DVMRP routing table.

Table 31-4 Commands for Clearing Caches, Tables, and Databases (continued)

Command	Purpose
clear ip igmp group [<i>group-name</i> <i>group-address</i> <i>interface</i>]	Delete entries from the IGMP cache.
clear ip mroute { <i>*</i> <i>group</i> [<i>source</i>]}	Delete entries from the IP multicast routing table.
clear ip pim auto-rp <i>rp-address</i>	Clear the Auto-RP cache.
clear ip sdr [<i>group-address</i> “ <i>session-name</i> ”]	Delete the Session Directory Protocol Version 2 cache or an sdr cache entry.

Displaying System and Network Statistics

You can display specific statistics, such as the contents of IP routing tables, caches, and databases.



Note

This release does not support per-route statistics.

You can display information to learn resource utilization and solve network problems. You can also display information about node reachability and discover the routing path your device’s packets are taking through the network.

You can use any of the privileged EXEC commands in [Table 31-5](#) to display various routing statistics:

Table 31-5 Commands for Displaying System and Network Statistics

Command	Purpose
ping [<i>group-name</i> <i>group-address</i>]	Send an ICMP Echo Request to a multicast group address.
show ip dvmrp route [<i>ip-address</i>]	Display the entries in the DVMRP routing table.
show ip igmp groups [<i>group-name</i> <i>group-address</i> <i>type number</i>]	Display the multicast groups that are directly connected to the switch and that were learned through IGMP.
show ip igmp interface [<i>type number</i>]	Display multicast-related information about an interface.
show ip mcache [<i>group</i> [<i>source</i>]]	Display the contents of the IP fast-switching cache.
show ip mpacket [<i>source-address</i> <i>name</i>] [<i>group-address</i> <i>name</i>] [detail]	Display the contents of the circular cache-header buffer.
show ip mroute [<i>group-name</i> <i>group-address</i>] [<i>source</i>] [summary] [count] [active kbps]	Display the contents of the IP multicast routing table.
show ip pim interface [<i>type number</i>] [count]	Display information about interfaces configured for PIM.
show ip pim neighbor [<i>type number</i>]	List the PIM neighbors discovered by the switch.
show ip pim rp [<i>group-name</i> <i>group-address</i>]	Display the RP routers associated with a sparse-mode multicast group.

Table 31-5 Commands for Displaying System and Network Statistics (continued)

Command	Purpose
show ip rpf { <i>source-address</i> <i>name</i> }	Display how the switch is doing Reverse-Path Forwarding (that is, from the unicast routing table, DVMRP routing table, or static mroutes).
show ip sdr [<i>group</i> “ <i>session-name</i> ” detail]	Display the Session Directory Protocol Version 2 cache.

Monitoring IP Multicast Routing

You can use the privileged EXEC commands in [Table 31-6](#) to monitor IP multicast routers, packets, and paths:

Table 31-6 Commands for Monitoring IP Multicast Routing

Command	Purpose
mrinfo [<i>hostname</i> <i>address</i>] [<i>source-address</i> <i>interface</i>]	Query a multicast router or multilayer switch about which neighboring multicast devices are peering with it.
mstat <i>source</i> [<i>destination</i>] [<i>group</i>]	Display IP multicast packet rate and loss information.
mtrace <i>source</i> [<i>destination</i>] [<i>group</i>]	Trace the path from a source to a destination branch for a multicast distribution tree for a given group.



Configuring MSDP

This chapter describes how to configure the Multicast Source Discovery Protocol (MSDP) on the Catalyst 3750 Metro switch. The MSDP connects multiple Protocol-Independent Multicast sparse-mode (PIM-SM) domains.

MSDP is not fully supported in this software release because of a lack of support for Multicast Border Gateway Protocol (MBGP), which works closely with MSDP. However, it is possible to create default peers that MSDP can operate with if MBGP is not running.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

This chapter consists of these sections:

- [Understanding MSDP, page 32-1](#)
- [Configuring MSDP, page 32-4](#)
- [Monitoring and Maintaining MSDP, page 32-19](#)

Understanding MSDP

MSDP allows multicast sources for a group to be known to all rendezvous points (RPs) in different domains. Each PIM-SM domain uses its own RPs and does not depend on RPs in other domains. An RP runs MSDP over the TCP to discover multicast sources in other domains.

An RP in a PIM-SM domain has an MSDP peering relationship with MSDP-enabled devices in another domain. The peering relationship occurs over a TCP connection, primarily exchanging a list of sources sending to multicast groups. The TCP connections between RPs are achieved by the underlying routing system. The receiving RP uses the source lists to establish a source path.

The purpose of this topology is to have domains discover multicast sources in other domains. If the multicast sources are of interest to a domain that has receivers, multicast data is delivered over the normal, source-tree building mechanism in PIM-SM. MSDP is also used to announce sources sending to a group. These announcements must originate at the domain's RP.

MSDP depends heavily on the Border Gateway Protocol (BGP) or MBGP for interdomain operation. We recommend that you run MSDP in RPs in your domain that are RPs for sources sending to global groups to be announced to the Internet.

MSDP Operation

Figure 32-1 shows MSDP operating between two MSDP peers. PIM uses MSDP as the standard mechanism to register a source with the RP of a domain. When MSDP is configured, this sequence occurs.


When a source sends its first multicast packet, the first-hop router (*designated router* or RP) directly connected to the source sends a PIM register message to the RP. The RP uses the register message to register the active source and to forward the multicast packet down the shared tree in the local domain. With MSDP configured, the RP also forwards a source-active (SA) message to all MSDP peers. The SA message identifies the source, the group the source is sending to, and the address of the RP or the originator ID (the IP address of the interface used as the RP address), if configured.

Each MSDP peer receives and forwards the SA message away from the originating RP to achieve peer reverse-path flooding (RPF). The MSDP device examines the BGP or MBGP routing table to discover which peer is the next hop toward the originating RP of the SA message. Such a peer is called an *RPF peer* (reverse-path forwarding peer). The MSDP device forwards the message to all MSDP peers other than the RPF peer. For information on how to configure an MSDP peer when BGP and MBGP are not supported, see the [“Configuring a Default MSDP Peer” section on page 32-4](#).

If the MSDP peer receives the same SA message from a non-RPF peer toward the originating RP, it drops the message. Otherwise, it forwards the message to all its MSDP peers.

The RP for a domain receives the SA message from an MSDP peer. If the RP has any join requests for the group the SA message describes and if the (*,G) entry exists with a nonempty outgoing interface list, the domain is interested in the group, and the RP triggers an (S,G) join toward the source. After the (S,G) join reaches the source’s DR, a branch of the source tree has been built from the source to the RP in the remote domain. Multicast traffic can now flow from the source across the source tree to the RP and then down the shared tree in the remote domain to the receiver.

Figure 32-1 MSDP Running Between RP Peers



MSDP Benefits

MSDP has these benefits:

- It breaks up the shared multicast distribution tree. You can make the shared tree local to your domain. Your local members join the local tree, and join messages for the shared tree never need to leave your domain.
- PIM sparse-mode domains can rely only on their own RPs, decreasing reliance on RPs in another domain. This increases security because you can prevent your sources from being known outside your domain.
- Domains with only receivers can receive data without globally advertising group membership.
- Global source multicast routing table state is not required, saving memory.

Configuring MSDP

These sections describe how to configure MSDP:

- [Default MSDP Configuration, page 32-4](#)
- [Configuring a Default MSDP Peer, page 32-4](#) (required)
- [Caching Source-Active State, page 32-6](#) (optional)
- [Requesting Source Information from an MSDP Peer, page 32-8](#) (optional)
- [Controlling Source Information that Your Switch Originates, page 32-8](#) (optional)
- [Controlling Source Information that Your Switch Forwards, page 32-12](#) (optional)
- [Controlling Source Information that Your Switch Receives, page 32-14](#) (optional)
- [Configuring an MSDP Mesh Group, page 32-16](#) (optional)
- [Shutting Down an MSDP Peer, page 32-16](#) (optional)
- [Including a Bordering PIM Dense-Mode Region in MSDP, page 32-17](#) (optional)
- [Configuring an Originating Address other than the RP Address, page 32-18](#) (optional)

Default MSDP Configuration

MSDP is not enabled, and no default MSDP peer exists.

Configuring a Default MSDP Peer


In this software release, because BGP and MBGP are not supported, you cannot configure an MSDP peer on the local switch by using the **ip msdp peer** global configuration command. Instead, you define a default MSDP peer (by using the **ip msdp default-peer** global configuration command) from which to accept all SA messages for the switch. The default MSDP peer must be a previously configured MSDP peer. Configure a default MSDP peer when the switch is not BGP- or MBGP-peering with an MSDP peer. If a single MSDP peer is configured, the switch always accepts all SA messages from that peer.

[Figure 32-2](#) shows a network in which default MSDP peers might be used. In [Figure 32-2](#), a customer who owns Switch B is connected to the Internet through two Internet service providers (ISPs), one owning Router A and the other owning Router C. They are not running BGP or MBGP between them. To learn about sources in the ISP's domain or in other domains, Switch B at the customer site identifies Router A as its default MSDP peer. Switch B advertises SA messages to both Router A and Router C but accepts SA messages only from Router A or only from Router C. If Router A is first in the configuration file, it is used if it is running. If Router A is not running, only then does Switch B accept SA messages from Router C. This is the default behavior without a prefix list.

If you specify a prefix list, the peer is a default peer only for the prefixes in the list. You can have multiple active default peers when you have a prefix list associated with each. When you do not have any prefix lists, you can configure multiple default peers, but only the first one is the active default peer as long as the router has connectivity to this peer and the peer is alive. If the first configured peer fails or the connectivity to this peer fails, the second configured peer becomes the active default, and so on.

The ISP probably uses a prefix list to define which prefixes it accepts from the customer's router.

Figure 32-2 Default MSDP Peer Network



Beginning in privileged EXEC mode, follow these steps to specify a default MSDP peer. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip msdp default-peer <i>ip-address</i> <i>name</i> [prefix-list <i>list</i>]	Define a default peer from which to accept all MSDP SA messages. - For <i>ip-address</i> <i>name</i>, enter the IP address or Domain Name System (DNS) server name of the MSDP default peer. (Optional) For prefix-list <i>list</i>, enter the list name that specifies the peer to be the default peer only for the listed prefixes. You can have multiple active default peers when you have a prefix list associated with each. When you enter multiple ip msdp default-peer commands with the prefix-list keyword, you use all the default peers at the same time for different RP prefixes. This syntax is typically used in a service provider cloud that connects stub site clouds. When you enter multiple ip msdp default-peer commands without the prefix-list keyword, a single active peer accepts all SA messages. If that peer fails, the next configured default peer accepts all SA messages. This syntax is typically used at a stub site.

	Command	Purpose
Step 3	ip prefix-list <i>name</i> [description <i>string</i>] seq <i>number</i> { permit deny } <i>network</i> <i>length</i>	(Optional) Create a prefix list using the name specified in Step 2. (Optional) For description <i>string</i>, enter a description of up to 80 characters to describe this prefix list. - For seq <i>number</i>, enter the sequence number of the entry. The range is 1 to 4294967294. - The deny keyword denies access to matching conditions. - The permit keyword permits access to matching conditions. - For <i>network length</i>, specify the network number and length (in bits) of the network mask that is permitted or denied.
Step 4	ip msdp description { <i>peer-name</i> <i>peer-address</i> } <i>text</i>	(Optional) Configure a description for the specified peer to make it easier to identify in a configuration or in show command output. By default, no description is associated with an MSDP peer.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the default peer, use the **no ip msdp default-peer** *ip-address* | *name* global configuration command.

This example shows a partial configuration of Router A and Router C in Figure 32-2. Each of these ISPs have more than one customer (like the customer in Figure 32-2) who use default peering (no BGP or MBGP). In that case, they might have similar configurations. That is, they accept SAs only from a default peer if the SA is permitted by the corresponding prefix list.

Router A

```
Router(config)# ip msdp default-peer 10.1.1.1
Router(config)# ip msdp default-peer 10.1.1.1 prefix-list site-a
Router(config)# ip prefix-list site-b permit 10.0.0.0/8
```

Router C

```
Router(config)# ip msdp default-peer 10.1.1.1 prefix-list site-a
Router(config)# ip prefix-list site-b permit 10.0.0.0/8
```

Caching Source-Active State

By default, the switch does not cache source/group pairs from received SA messages. When the switch forwards the MSDP SA information, it does not store it in memory. Therefore, if a member joins a group soon after a SA message is received by the local RP, that member needs to wait until the next SA message to hear about the source. This delay is known as join latency.

If you want to sacrifice some memory in exchange for reducing the latency of the source information, you can configure the switch to cache SA messages.

Beginning in privileged EXEC mode, follow these steps to enable the caching of source/group pairs. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip msdp cache-sa-state [list <i>access-list-number</i>]	Enable the caching of source/group pairs (create an SA state). Those pairs that pass the access list are cached. For list <i>access-list-number</i> , the range is 100 to 199.
Step 3	access-list <i>access-list-number</i> { deny permit } <i>protocol source source-wildcard destination destination-wildcard</i>	Create an IP extended access list, repeating the command as many times as necessary. - For <i>access-list-number</i>, the range is 100 to 199. Enter the same number created 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>protocol</i>, enter ip as the protocol name. - For <i>source</i>, enter the number of the network or host from which the packet is being sent. - 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. - For <i>destination</i>, enter the number of the network or host to which the packet is being sent. - For <i>destination-wildcard</i>, enter the wildcard bits in dotted decimal notation to be applied to the destination. 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.



Note

An alternative to this command is the **ip msdp sa-request** global configuration command, which causes the switch to send an SA request message to the MSDP peer when a new member for a group becomes active. For more information, see the next section.

To return to the default setting (no SA state is created), use the **no ip msdp cache-sa-state** global configuration command.

This example shows how to enable the cache state for all sources in 171.69.0.0/16 sending to groups 224.2.0.0/16:

```
Switch(config)# ip msdp cache-sa-state 100
Switch(config)# access-list 100 permit ip 171.69.0.0 0.0.255.255 224.2.0.0 0.0.255.255
```

Requesting Source Information from an MSDP Peer

Local RPs can send SA requests and get immediate responses for all active sources for a given group. By default, the switch does not send any SA request messages to its MSDP peers when a new member joins a group and wants to receive multicast traffic. The new member waits to receive the next periodic SA message.

If you want a new member of a group to learn the active multicast sources in a connected PIM sparse-mode domain that are sending to a group, configure the switch to send SA request messages to the specified MSDP peer when a new member joins a group. The peer replies with the information in its SA cache. If the peer does not have a cache configured, this command has no result. Configuring this feature reduces join latency but sacrifices memory.

Beginning in privileged EXEC mode, follow these steps to configure the switch to send SA request messages to the MSDP peer when a new member joins a group and wants to receive multicast traffic. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip msdp sa-request { <i>ip-address</i> <i>name</i> }	Configure the switch to send SA request messages to the specified MSDP peer. For <i>ip-address</i> <i>name</i> , enter the IP address or name of the MSDP peer from which the local switch requests SA messages when a new member for a group becomes active. Repeat the command for each MSDP peer that you want to supply with SA messages.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no ip msdp sa-request** {*ip-address* | *name*} global configuration command.

This example shows how to configure the switch to send SA request messages to the MSDP peer at 171.69.1.1:

```
Switch(config)# ip msdp sa-request 171.69.1.1
```

Controlling Source Information that Your Switch Originates

You can control the multicast source information that originates with your switch:

- Sources you advertise (based on your sources)
- Receivers of source information (based on knowing the requestor)

For more information, see the [“Redistributing Sources” section on page 32-9](#) and the [“Filtering Source-Active Request Messages” section on page 32-11](#).

Redistributing Sources

SA messages originate on RPs to which sources have registered. By default, any source that registers with an RP is advertised. The *A flag* is set in the RP when a source is registered, which means the source is advertised in an SA unless it is filtered.

Beginning in privileged EXEC mode, follow these steps to further restrict which registered sources are advertised. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip msdp redistribute [list <i>access-list-name</i>] [asn <i>aspath-access-list-number</i>] [route-map <i>map</i>]	Configure which (S,G) entries from the multicast routing table are advertised in SA messages. By default, only sources within the local domain are advertised. • (Optional) For list <i>access-list-name</i>, enter the name or number of an IP standard or extended access list. The range is 1 to 99 for standard access lists and 100 to 199 for extended lists. The access list controls which local sources are advertised and to which groups they send. • (Optional) For asn <i>aspath-access-list-number</i>, enter the IP standard or extended access list number in the range 1 to 199. This access list number must also be configured in the ip as-path access-list command. • (Optional) For route-map <i>map</i>, enter the IP standard or extended access list number in the range 1 to 199. This access list number must also be configured in the ip as-path access-list command. The switch advertises (S,G) pairs according to the access list or autonomous system path access list.

	Command	Purpose
Step 3	access-list <i>access-list-number</i> {deny permit} <i>source</i> [<i>source-wildcard</i>] or access-list <i>access-list-number</i> {deny permit} <i>protocol source source-wildcard destination destination-wildcard</i>	Create an IP standard access list, repeating the command as many times as necessary. or Create an IP extended access list, repeating the command as many times as necessary. - For <i>access-list-number</i>, the range is 1 to 99 for standard access lists and 100 to 199 for extended lists. Enter the same number created 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>protocol</i>, enter ip as the protocol name. - For <i>source</i>, enter the number of the network or host from which the packet is being sent. - 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. - For <i>destination</i>, enter the number of the network or host to which the packet is being sent. - For <i>destination-wildcard</i>, enter the wildcard bits in dotted decimal notation to be applied to the destination. 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the filter, use the **no ip msdp redistribute** global configuration command.

Filtering Source-Active Request Messages

By default, only switches that are caching SA information can respond to SA requests. By default, such a switch honors all SA request messages from its MSDP peers and supplies the IP addresses of the active sources.

However, you can configure the switch to ignore all SA requests from an MSDP peer. You can also honor only those SA request messages from a peer for groups described by a standard access list. If the groups in the access list pass, SA request messages are accepted. All other such messages from the peer for other groups are ignored.

Beginning in privileged EXEC mode, follow these steps to configure one of these options. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip msdp filter-sa-request <i>ip-address</i> <i>name</i> or ip msdp filter-sa-request { <i>ip-address</i> <i>name</i> } list <i>access-list-number</i>	Filter all SA request messages from the specified MSDP peer. or Filter SA request messages from the specified MSDP peer for groups that pass the standard access list. The access list describes a multicast group address. The range for the access-list-number is 1 to 99.
Step 3	access-list <i>access-list-number</i> { deny permit } <i>source</i> [<i>source-wildcard</i>]	Create an IP standard access list, repeating the command as many times as necessary. - For <i>access-list-number</i>, the range is 1 to 99. - 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 number of the network or host from which the packet is being sent. (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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no ip msdp filter-sa-request** {*ip-address* | *name*} global configuration command.

This example shows how to configure the switch to filter SA request messages from the MSDP peer at 171.69.2.2. SA request messages from sources on network 192.4.22.0 pass access list 1 and are accepted; all others are ignored.

```
Switch(config)# ip msdp filter sa-request 171.69.2.2 list 1
Switch(config)# access-list 1 permit 192.4.22.0 0.0.0.255
```

Controlling Source Information that Your Switch Forwards

By default, the switch forwards all SA messages it receives to all its MSDP peers. However, you can prevent outgoing messages from being forwarded to a peer by using a filter or by setting a time-to-live (TTL) value. These methods are described in the next sections.

Using a Filter

By creating a filter, you can perform one of these actions:

- Filter all source/group pairs
- Specify an IP extended access list to pass only certain source/group pairs
- Filter based on match criteria in a route map

Beginning in privileged EXEC mode, follow these steps to apply a filter. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip msdp sa-filter out <i>ip-address</i> <i>name</i>	Filter all SA messages to the specified MSDP peer.
	or	or
	ip msdp sa-filter out { <i>ip-address</i> <i>name</i> } list <i>access-list-number</i>	To the specified peer, pass only those SA messages that pass the IP extended access list. The range for the extended <i>access-list-number</i> is 100 to 199. If both the list and the route-map keywords are used, all conditions must be true to pass any (S,G) pair in outgoing SA messages.
	or	or
	ip msdp sa-filter out { <i>ip-address</i> <i>name</i> } route-map <i>map-tag</i>	To the specified MSDP peer, pass only those SA messages that meet the match criteria in the route map <i>map-tag</i> . If all match criteria are true, a permit from the route map passes routes through the filter. A deny filters routes.

	Command	Purpose
Step 3	access-list <i>access-list-number</i> { deny permit } <i>protocol source source-wildcard destination destination-wildcard</i>	(Optional) Create an IP extended access list, repeating the command as many times as necessary. - For <i>access-list-number</i>, enter the 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>protocol</i>, enter ip as the protocol name. - For <i>source</i>, enter the number of the network or host from which the packet is being sent. - 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. - For <i>destination</i>, enter the number of the network or host to which the packet is being sent. - For <i>destination-wildcard</i>, enter the wildcard bits in dotted decimal notation to be applied to the destination. 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the filter, use the **no ip msdp sa-filter out** {*ip-address* | *name*} [**list** *access-list-number*] [**route-map** *map-tag*] global configuration command.

This example shows how to allow only (S,G) pairs that pass access list 100 to be forwarded in an SA message to the peer named *switch.cisco.com*:

```
Switch(config)# ip msdp peer switch.cisco.com connect-source gigabitethernet1/0/1
Switch(config)# ip msdp sa-filter out switch.cisco.com list 100
Switch(config)# access-list 100 permit ip 171.69.0.0 0.0.255.255 224.20.0 0.0.255.255
```

Using TTL to Limit the Multicast Data Sent in SA Messages

You can use a TTL value to control what data is encapsulated in the first SA message for every source. Only multicast packets with an IP-header TTL greater than or equal to the *ttl* argument are sent to the specified MSDP peer. For example, you can limit internal traffic to a TTL of 8. If you want other groups to go to external locations, you must send those packets with a TTL greater than 8.

Beginning in privileged EXEC mode, follow these steps to establish a TTL threshold. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip msdp ttl-threshold { <i>ip-address</i> <i>name</i> } <i>ttl</i>	Limit which multicast data is encapsulated in the first SA message to the specified MSDP peer. - For <i>ip-address</i> <i>name</i>, enter the IP address or name of the MSDP peer to which the TTL limitation applies. - For <i>ttl</i>, enter the TTL value. The default is 0, which means all multicast data packets are forwarded to the peer until the TTL is exhausted. The range is 0 to 255.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no ip msdp ttl-threshold** {*ip-address* | *name*} global configuration command.

Controlling Source Information that Your Switch Receives

By default, the switch receives all SA messages that its MSDP RPF peers send to it. However, you can control the source information that you receive from MSDP peers by filtering incoming SA messages. In other words, you can configure the switch to not accept them.

You can perform one of these actions:

- Filter all incoming SA messages from an MSDP peer
- Specify an IP extended access list to pass certain source/group pairs
- Filter based on match criteria in a route map

Beginning in privileged EXEC mode, follow these steps to apply a filter. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip msdp sa-filter in <i>ip-address name</i> or ip msdp sa-filter in { <i>ip-address name</i> } list <i>access-list-number</i> or ip msdp sa-filter in { <i>ip-address name</i> } route-map <i>map-tag</i>	Filter all SA messages from the specified MSDP peer. or From the specified peer, pass only those SA messages that pass the IP extended access list. The range for the extended <i>access-list-number</i> is 100 to 199. If both the list and the route-map keywords are used, all conditions must be true to pass any (S,G) pair in incoming SA messages. or From the specified MSDP peer, pass only those SA messages that meet the match criteria in the route map <i>map-tag</i> . If all match criteria are true, a permit from the route map passes routes through the filter. A deny will filter routes.
Step 3	access-list <i>access-list-number</i> { deny permit } <i>protocol source source-wildcard destination destination-wildcard</i>	(Optional) Create an IP extended access list, repeating the command as many times as necessary. - For <i>access-list-number</i>, enter the 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>protocol</i>, enter ip as the protocol name. - For <i>source</i>, enter the number of the network or host from which the packet is being sent. - 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. - For <i>destination</i>, enter the number of the network or host to which the packet is being sent. - For <i>destination-wildcard</i>, enter the wildcard bits in dotted decimal notation to be applied to the destination. 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.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the filter, use the **no ip msdp sa-filter in** {*ip-address | name*} [**list** *access-list-number*] [**route-map** *map-tag*] global configuration command.

This example shows how to filter all SA messages from the peer named *switch.cisco.com*:

```
Switch(config)# ip msdp peer switch.cisco.com connect-source gigabitethernet1/0/1
Switch(config)# ip msdp sa-filter in switch.cisco.com
```

Configuring an MSDP Mesh Group

An MSDP mesh group is a group of MSDP speakers that have fully meshed MSDP connectivity among one another. Any SA messages received from a peer in a mesh group are not forwarded to other peers in the same mesh group. Thus, you reduce SA message flooding and simplify peer-RPF flooding. Use the **ip msdp mesh-group** global configuration command when there are multiple RPs within a domain. It is especially used to send SA messages across a domain. You can configure multiple mesh groups (with different names) in a single switch.

Beginning in privileged EXEC mode, follow these steps to create a mesh group. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip msdp mesh-group <i>name</i> { <i>ip-address</i> <i>name</i> }	Configure an MSDP mesh group, and specify the MSDP peer belonging to that mesh group. By default, the MSDP peers do not belong to a mesh group. - For <i>name</i>, enter the name of the mesh group. - For <i>ip-address</i> <i>name</i>, enter the IP address or name of the MSDP peer to be a member of the mesh group.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.
Step 6		Repeat this procedure on each MSDP peer in the group.

To remove an MSDP peer from a mesh group, use the **no ip msdp mesh-group** *name* {*ip-address* | *name*} global configuration command.

Shutting Down an MSDP Peer

If you want to configure many MSDP commands for the same peer and you do not want the peer to become active, you can shut down the peer, configure it, and later bring it up. When a peer is shut down, the TCP connection is terminated and is not restarted. You can also shut down an MSDP session without losing configuration information for the peer.

Beginning in privileged EXEC mode, follow these steps to shut down a peer. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip msdp shutdown { <i>peer-name</i> <i>peer address</i> }	Administratively shut down the specified MSDP peer without losing configuration information. For <i>peer-name</i> <i>peer address</i> , enter the IP address or name of the MSDP peer to shut down.
Step 3	end	Return to privileged EXEC mode.

	Command	Purpose
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To bring the peer back up, use the **no ip msdp shutdown** {*peer-name* | *peer address*} global configuration command. The TCP connection is reestablished

Including a Bordering PIM Dense-Mode Region in MSDP

You can configure MSDP on a switch that borders a PIM sparse-mode region with a dense-mode region. By default, active sources in the dense-mode region do not participate in MSDP.



Note

We do not recommend using the **ip msdp border sa-address** global configuration command. It is better to configure the border router in the sparse-mode domain to proxy-register sources in the dense-mode domain to the RP of the sparse-mode domain and have the sparse-mode domain use standard MSDP procedures to advertise these sources.

Beginning in privileged EXEC mode, follow these steps to configure the border router to send SA messages for sources active in the dense-mode region to the MSDP peers. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip msdp border sa-address <i>interface-id</i>	Configure the switch on the border between a dense-mode and sparse-mode region to send SA messages about active sources in the dense-mode region. For <i>interface-id</i> , specify the interface from which the IP address is derived and used as the RP address in SA messages. The IP address of the interface is used as the Originator-ID, which is the RP field in the SA message.
Step 3	ip msdp redistribute [<i>list access-list-name</i>] [<i>asn aspath-access-list-number</i>] [<i>route-map map</i>]	Configure which (S,G) entries from the multicast routing table are advertised in SA messages. For more information, see the “Redistributing Sources” section on page 32-9 .
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Note that the **ip msdp originator-id** global configuration command also identifies an interface to be used as the RP address. If you configure both the **ip msdp border sa-address** and the **ip msdp originator-id** global configuration commands, the address derived from the **ip msdp originator-id** command determines the RP address.

To return to the default setting (active sources in the dense-mode region do not participate in MSDP), use the **no ip msdp border sa-address** *interface-id* global configuration command.

Configuring an Originating Address other than the RP Address

You can allow an MSDP speaker that originates an SA message to use the IP address of the interface as the RP address in the SA message by changing the Originator ID. You might change the Originator ID in one of these cases:

- If you configure a logical RP on multiple switches in an MSDP mesh group.
- If you have a switch that borders a PIM sparse-mode domain and a dense-mode domain. If a switch borders a dense-mode domain for a site, and sparse-mode is being used externally, you might want dense-mode sources to be known to the outside world. Because this switch is not an RP, it would not have an RP address to use in an SA message. Therefore, this command provides the RP address by specifying the address of the interface.

Beginning in privileged EXEC mode, follow these steps to allow an MSDP speaker that originates an SA message to use the IP address on the interface as the RP address in the SA message. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip msdp originator-id <i>interface-id</i>	Configures the RP address in SA messages to be the address of the originating device interface. For <i>interface-id</i> , specify the interface on the local switch.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

If you configure both the **ip msdp border sa-address** and the **ip msdp originator-id** global configuration commands, the address derived from the **ip msdp originator-id** command determines the address of the RP.

To prevent the RP address from being derived in this way, use the **no ip msdp originator-id** *interface-id* global configuration command.

Monitoring and Maintaining MSDP

To monitor MSDP SA messages, peers, state, or peer status, use one or more of the privileged EXEC commands in [Table 32-1](#):

Table 32-1 Commands for Monitoring and Maintaining MSDP

Command	Purpose
debug ip msdp [<i>peer-address</i> <i>name</i>] [detail] [routes]	Debugs an MSDP activity.
debug ip msdp resets	Debugs MSDP peer reset reasons.
show ip msdp count [<i>autonomous-system-number</i>]	Displays the number of sources and groups originated in SA messages from each autonomous system. The ip msdp cache-sa-state command must be configured for this command to produce any output.
show ip msdp peer [<i>peer-address</i> <i>name</i>]	Displays detailed information about an MSDP peer.
show ip msdp sa-cache [<i>group-address</i> <i>source-address</i> <i>group-name</i> <i>source-name</i>] [<i>autonomous-system-number</i>]	Displays (S,G) state learned from MSDP peers.
show ip msdp summary	Displays MSDP peer status and SA message counts.

To clear MSDP connections, statistics, or SA cache entries, use the privileged EXEC commands in [Table 32-2](#):

Table 32-2 Commands for Clearing MSDP Connections, Statistics, or SA Cache Entries

Command	Purpose
clear ip msdp peer <i>peer-address</i> <i>name</i>	Clears the TCP connection to the specified MSDP peer, resetting all MSDP message counters.
clear ip msdp statistics [<i>peer-address</i> <i>name</i>]	Clears statistics counters for one or all the MSDP peers without resetting the sessions.
clear ip msdp sa-cache [<i>group-address</i> <i>name</i>]	Clears the SA cache entries for all entries, all sources for a specific group, or all entries for a specific source/group pair.



Configuring Fallback Bridging

This chapter describes how to configure fallback bridging (VLAN bridging) on the Catalyst 3750 Metro switch. With fallback bridging, you can forward non-IP packets that the switch does not route between VLAN bridge domains and routed ports.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco IOS Bridging and IBM Networking Command Reference for Release 12.1*.

This chapter consists of these sections:

- [Understanding Fallback Bridging, page 33-1](#)
- [Configuring Fallback Bridging, page 33-2](#)
- [Monitoring and Maintaining Fallback Bridging, page 33-10](#)

Understanding Fallback Bridging

With fallback bridging, the switch bridges together two or more VLANs or routed ports, essentially connecting multiple VLANs within one bridge domain. Fallback bridging forwards traffic that the switch does not route and forwards traffic belonging to a nonroutable protocol such as DECnet.

A VLAN bridge domain is represented with switch virtual interfaces (SVIs). A set of SVIs and routed ports (which do not have any VLANs associated with them) can be configured (grouped together) to form a bridge group. Recall that an SVI represents a VLAN of switch ports as one interface to the routing or bridging function in the system. You associate only one SVI with a VLAN, and you configure an SVI for a VLAN only when you want to route between VLANs, to fallback-bridge nonroutable protocols between VLANs, or to provide IP host connectivity to the switch. A routed port is a physical port that acts like a port on a router, but it is not connected to a router. A routed port is not associated with a particular VLAN, does not support VLAN subinterfaces, but behaves like a normal routed port. For more information about SVIs and routed ports, see [Chapter 9, “Configuring Interface Characteristics.”](#)

A bridge group is an internal organization of network interfaces on a switch. You cannot use bridge groups to identify traffic switched within the bridge group outside the switch on which they are defined. Bridge groups on the switch function as distinct bridges; that is, bridged traffic and bridge protocol data units (BPDUs) are not exchanged between different bridge groups on a switch.

Fallback bridging does not allow the spanning trees from the VLANs being bridged to collapse. Each VLAN has its own spanning-tree instance and a separate spanning tree, called the VLAN-bridge spanning tree, which runs on top of the bridge group to prevent loops.


The switch creates a VLAN-bridge spanning-tree instance when a bridge group is created. The switch runs the bridge group and treats the SVIs and routed ports in the bridge group as its spanning-tree ports.

These are the reasons for placing network interfaces into a bridge group:

- To bridge all nonrouted traffic among the network interfaces making up the bridge group. If the packet destination address is in the bridge table, the packet is forwarded on a single interface in the bridge group. If the packet destination address is not in the bridge table, the packet is flooded on all forwarding interfaces in the bridge group. A source MAC address is learned on a bridge group only when the address is learned on a VLAN (the reverse is not true).
- To participate in the spanning-tree algorithm by receiving, and in some cases sending, BPDUs on the LANs to which they are attached. A separate spanning-tree process runs for each configured bridge group. Each bridge group participates in a separate spanning-tree instance. A bridge group establishes a spanning-tree instance based on the BPDUs it receives on only its member ports. If the bridge STP BPDU is received on a port whose VLAN does not belong to a bridge group, the BPDU is flooded on all the forwarding ports of the VLAN.

Figure 33-1 shows a fallback bridging network example. The switch has two ports configured as SVIs with different assigned IP addresses and attached to two different VLANs. Another port is configured as a routed port with its own IP address. If all three of these ports are assigned to the same bridge group, non-IP protocol frames can be forwarded among the end stations connected to the switch even though they are on different networks and in different VLANs. IP addresses do not need to be assigned to routed ports or SVIs for fallback bridging to work.

Figure 33-1 Fallback Bridging Network Example



Configuring Fallback Bridging

These sections describe how to configure fallback bridging on your switch:

- [Default Fallback Bridging Configuration, page 33-3](#)
- [Fallback Bridging Configuration Guidelines, page 33-3](#)
- [Creating a Bridge Group, page 33-3](#) (required)
- [Adjusting Spanning-Tree Parameters, page 33-5](#) (optional)

Default Fallback Bridging Configuration

Table 33-1 shows the default fallback bridging configuration.

Table 33-1 Default Fallback Bridging Configuration

Feature	Default Setting
Bridge groups	None are defined or assigned to a port. No VLAN-bridge STP is defined.
Switch forwards frames for stations that it has dynamically learned	Enabled.
Spanning tree parameters:	
• Switch priority	• 32768.
• Port priority	• 128.
• Port path cost	• 10 Mbps: 100. 100 Mbps: 19. 1000 Mbps: 4.
• Hello BPDU interval	• 2 seconds.
• Forward-delay interval	• 20 seconds.
• Maximum idle interval	• 30 seconds.

Fallback Bridging Configuration Guidelines

Up to 32 bridge groups can be configured on the switch.

An interface (an SVI or routed port) can be a member of only one bridge group.

Use a bridge group for each separately bridged (topologically distinct) network connected to the switch.

Creating a Bridge Group

To configure fallback bridging for a set of SVIs or routed ports, these interfaces must be assigned to bridge groups. All interfaces in the same group belong to the same bridge domain. Each SVI or routed port can be assigned to only one bridge group.



Note

The protected port feature is not compatible with fallback bridging. When fallback bridging is enabled, it is possible for packets to be forwarded from one protected port on a switch to another protected port on the same switch if the ports are in different VLANs.

Beginning in privileged EXEC mode, follow these steps to create a bridge group and to assign an interface to it. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	bridge <i>bridge-group</i> protocol vlan-bridge	Assign a bridge group number, and specify the VLAN-bridge spanning-tree protocol to run in the bridge group. The ibm and dec keywords are not supported. For <i>bridge-group</i> , specify the bridge group number. The range is 1 to 255. You can create up to 32 bridge groups. Frames are bridged only among interfaces in the same group.
Step 3	interface <i>interface-id</i>	Specify the interface on which you want to assign the bridge group, and enter interface configuration mode. The specified interface must be one of these: - A routed port: a physical port that you have configured as a Layer 3 port by entering the no switchport interface configuration command. - An SVI: a VLAN interface that you created by using the interface vlan <i>vlan-id</i> global configuration command. Note You can assign an IP address to the routed port or to the SVI, but it is not required.
Step 4	bridge-group <i>bridge-group</i>	Assign the interface to the bridge group created in Step 2. By default, the interface is not assigned to any bridge group. An interface can be assigned to only one bridge group.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a bridge group, use the **no bridge *bridge-group*** global configuration command. The **no bridge *bridge-group*** command automatically removes all SVIs and routed ports from that bridge group. To remove an interface from a bridge group and to remove the bridge group, use the **no bridge-group *bridge-group*** interface configuration command.

This example shows how to create bridge group 10, to specify that the VLAN-bridge STP runs in the bridge group, to define a port as a routed port, and to assign the port to the bridge group:

```
Switch(config)# bridge 10 protocol vlan-bridge
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# no switchport
Switch(config-if)# no shutdown
Switch(config-if)# bridge-group 10
```

This example shows how to create bridge group 10 and to specify that the VLAN-bridge STP runs in the bridge group. It defines a port as an SVI and assigns this port to VLAN 2 and to the bridge group:

```
Switch(config)# bridge 10 protocol vlan-bridge
Switch(config)# vlan 2
Switch(config-vlan)# exit
Switch(config)# interface vlan 2
Switch(config-if)# bridge-group 10
Switch(config-if)# no shutdown
Switch(config-if)# exit
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan 2
Switch(config-if)# no shutdown
```

Adjusting Spanning-Tree Parameters

You might need to adjust certain spanning-tree parameters if the default values are not suitable. You configure parameters affecting the entire spanning tree by using variations of the **bridge** global configuration command. You configure port-specific parameters by using variations of the **bridge-group** interface configuration command.

You can adjust spanning-tree parameters by performing any of the tasks in these sections:

- [Changing the VLAN-Bridge Spanning-Tree Priority, page 33-5](#) (optional)
- [Changing the Interface Priority, page 33-6](#) (optional)
- [Assigning a Path Cost, page 33-7](#) (optional)
- [Adjusting BPDU Intervals, page 33-8](#) (optional)
- [Disabling the Spanning Tree on an Interface, page 33-10](#) (optional)



Note

Only network administrators with a good understanding of how switches and STP function should make adjustments to spanning-tree parameters. Poorly planned adjustments can have a negative impact on performance. A good source on switching is the IEEE 802.1D specification. For more information, refer to the “References and Recommended Reading” appendix in the *Cisco IOS Configuration Fundamentals Command Reference*.

Changing the VLAN-Bridge Spanning-Tree Priority

You can globally configure the VLAN-bridge spanning-tree priority of a switch when it ties with another switch for the position as the root switch. You also can configure the likelihood that the switch will be selected as the root switch.

Beginning in privileged EXEC mode, follow these steps to change the switch priority. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	bridge <i>bridge-group</i> <i>priority number</i>	Change the VLAN-bridge spanning-tree priority of the switch. - For <i>bridge-group</i>, specify the bridge group number. The range is 1 to 255. - For <i>number</i>, enter a number from 0 to 65535. The default is 32768. The lower the number, the more likely the switch will be chosen as the root.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To return to the default setting, use the **no bridge *bridge-group* *priority*** global configuration command. To change the priority on a port, use the **bridge-group *priority*** interface configuration command (described in the next section).

This example shows how to set the switch priority to 100 for bridge group 10:

```
Switch(config)# bridge 10 priority 100
```

Changing the Interface Priority

You can change the priority for a port. When two switches tie for position as the root switch, you configure a port priority to break the tie. The switch with the lowest port value is elected.

Beginning in privileged EXEC mode, follow these steps to change the port priority. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port to set the priority, and enter interface configuration mode.
Step 3	bridge-group <i>bridge-group</i> <i>priority number</i>	Change the priority of a port. - For <i>bridge-group</i>, specify the bridge group number. The range is 1 to 255. - For <i>number</i>, enter a number from 0 to 255 in increments of 4. The lower the number, the more likely that the port on the switch will be chosen as the root. The default is 128.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entry.
Step 6	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To return to the default setting, use the **no bridge-group *bridge-group* *priority*** interface configuration command.

This example shows how to change the priority to 20 on a port in bridge group 10:

```
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# bridge-group 10 priority 20
```

Assigning a Path Cost

Each port has a path cost associated with it. By convention, the path cost is 1000/data rate of the attached LAN, in Mbps.

Beginning in privileged EXEC mode, follow these steps to assign a path cost. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port to set the path cost, and enter interface configuration mode.
Step 3	bridge-group <i>bridge-group</i> path-cost <i>cost</i>	Assign the path cost of a port. - For <i>bridge-group</i>, specify the bridge group number. The range is 1 to 255. - For <i>cost</i>, enter a number from 0 to 65535. The higher the value, the higher the cost. - For 10 Mbps, the default path cost is 100. - For 100 Mbps, the default path cost is 19. - For 1000 Mbps, the default path cost is 4.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entry.
Step 6	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To return to the default path cost, use the **no bridge-group** *bridge-group* **path-cost** interface configuration command.

This example shows how to change the path cost to 20 on a port in bridge group 10:

```
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# bridge-group 10 path-cost 20
```

Adjusting BPDU Intervals

You can adjust BPDU intervals as described in these sections:

- [Adjusting the Interval between Hello BPDUs, page 33-8](#) (optional)
- [Changing the Forward-Delay Interval, page 33-9](#) (optional)
- [Changing the Maximum-Idle Interval, page 33-9](#) (optional)

**Note**

Each switch in a spanning tree adopts the interval between hello BPDUs, the forward delay interval, and the maximum idle interval parameters of the root switch, regardless of what its individual configuration might be.

Adjusting the Interval between Hello BPDUs

Beginning in privileged EXEC mode, follow these step to adjust the interval between hello BPDUs. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	bridge <i>bridge-group</i> hello-time <i>seconds</i>	Specify the interval between hello BPDUs. • For <i>bridge-group</i>, specify the bridge group number. The range is 1 to 255. • For <i>seconds</i>, enter a number from 1 to 10. The default is 2.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To return to the default setting, use the **no bridge *bridge-group* hello-time** global configuration command.

This example shows how to change the hello interval to 5 seconds in bridge group 10:

```
Switch(config)# bridge 10 hello-time 5
```

Changing the Forward-Delay Interval

The forward-delay interval is the amount of time spent listening for topology change information after a port has been activated for switching and before forwarding actually begins.

Beginning in privileged EXEC mode, follow these steps to change the forward-delay interval. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	bridge <i>bridge-group</i> forward-time <i>seconds</i>	Specify the forward-delay interval. - For <i>bridge-group</i>, specify the bridge group number. The range is 1 to 255. - For <i>seconds</i>, enter a number from 4 to 200. The default is 20.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To return to the default setting, use the **no bridge *bridge-group* forward-time** global configuration command.

This example shows how to change the forward-delay interval to 10 seconds in bridge group 10:

```
Switch(config)# bridge 10 forward-time 10
```

Changing the Maximum-Idle Interval

If a switch does not receive BPDUs from the root switch within a specified interval, it recomputes the spanning-tree topology.

Beginning in privileged EXEC mode, follow these steps to change the maximum-idle interval (maximum aging time). This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	bridge <i>bridge-group</i> max-age <i>seconds</i>	Specify the interval that the switch waits to hear BPDUs from the root switch. - For <i>bridge-group</i>, specify the bridge group number. The range is 1 to 255. - For <i>seconds</i>, enter a number from 6 to 200. The default is 30.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To return to the default setting, use the **no bridge *bridge-group* max-age** global configuration command.

This example shows how to change the maximum-idle interval to 30 seconds in bridge group 10:

```
Switch(config)# bridge 10 max-age 30
```

Disabling the Spanning Tree on an Interface

When a loop-free path exists between any two switched subnetworks, you can prevent BPDUs generated in one switching subnetwork from impacting devices in the other switching subnetwork, yet still permit switching throughout the network as a whole. For example, when switched LAN subnetworks are separated by a WAN, BPDUs can be prevented from traveling across the WAN link.

Beginning in privileged EXEC mode, follow these steps to disable spanning tree on a port. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port, and enter interface configuration mode.
Step 3	bridge-group <i>bridge-group</i> spanning-disabled	Disable spanning tree on the port. For <i>bridge-group</i> , specify the bridge group number. The range is 1 to 255.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entry.
Step 6	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To re-enable spanning tree on the port, use the **no bridge-group** *bridge-group* **spanning-disabled** interface configuration command.

This example shows how to disable spanning tree on a port in bridge group 10:

```
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# bridge group 10 spanning-disabled
```

Monitoring and Maintaining Fallback Bridging

To monitor and maintain the network, use one or more of the privileged EXEC commands in [Table 33-2](#):

Table 33-2 Commands for Monitoring and Maintaining Fallback Bridging

Command	Purpose
clear bridge <i>bridge-group</i>	Removes any learned entries from the forwarding database.
show bridge [<i>bridge-group</i>] group	Displays details about the bridge group.
show bridge [<i>bridge-group</i>] [<i>interface-id</i> <i>mac-address</i> verbose]	Displays MAC addresses learned in the bridge group.

For information about the fields in these displays, refer to the *Cisco IOS Bridging and IBM Networking Command Reference for Release 12.1*.



Troubleshooting

This chapter describes how to identify and resolve software problems related to the Cisco IOS software on the Catalyst 3750 Metro switch.

Additional troubleshooting information is provided in the hardware installation guide.



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 Command Summary for Release 12.1*.

This chapter consists of these sections:

- [Recovering from Corrupted Software By Using the XMODEM Protocol, page 34-2](#)
- [Recovering from a Lost or Forgotten Password, page 34-3](#)



Note Recovery procedures require that you have physical access to the switch.

- [Preventing Autonegotiation Mismatches, page 34-7](#)
- [SFP Module Security and Identification, page 34-8](#)
- [Using Ping, page 34-8](#)
- [Using Layer 2 Traceroute, page 34-10](#)
- [Using IP Traceroute, page 34-11](#)
- [Using Debug Commands, page 34-13](#)
- [Using the show platform forward Command, page 34-14](#)
- [Using the crashinfo File, page 34-17](#)

Recovering from Corrupted Software By Using the XMODEM Protocol

Switch software can be corrupted during an upgrade, by downloading the wrong file to the switch, and by deleting the image file. In all of these cases, the switch does not pass the power-on self-test (POST), and there is no connectivity.

This procedure uses the XMODEM Protocol to recover from a corrupt or wrong image file. There are many software packages that support the XMODEM Protocol, and this procedure is largely dependent on the emulation software you are using.

This recovery procedure requires that you have physical access to the switch.

Step 1 From your PC, download the software image tar file (*image_filename.tar*) from Cisco.com.

The Cisco IOS image is stored as a bin file in a directory in the tar file. For information about locating the software image files on Cisco.com, refer to the release notes.

Step 2 Extract the bin file from the tar file.

- If you are using Windows, use a zip program that is capable of reading a tar file. Use the zip program to navigate to and extract the bin file.
- If you are using UNIX, follow these steps:
 1. Display the contents of the tar file by using the **tar -tvf <image_filename.tar>** UNIX command.
 2. Locate the bin file name in the display and extract it by using the **tar -xvf <image_filename.tar> <image_filename.bin>** UNIX command.


```
switch% tar -xvf image_filename.tar image_filename.bin
```
 3. Verify that the bin file was extracted by using the **ls -l <image_filename.bin>** UNIX command. The bin file name (*image_filename.bin*) should appear in the output.


```
switch% ls -l image_filename.bin
```

Step 3 Connect your PC with terminal-emulation software supporting the XMODEM Protocol to the switch console port.

Step 4 Set the line speed on the emulation software to 9600 baud.

Step 5 Unplug the switch power cord.

Step 6 Press the **Mode** button, and at the same time, reconnect the power cord to the switch.

You can release the **Mode** button a second or two after the LED above port 1 goes off. Several lines of information about the software appear along with instructions:

The system has been interrupted prior to initializing the flash file system. The following commands will initialize the flash file system, and finish loading the operating system software#

```
flash_init
load_helper
boot
```

Step 7 Initialize the Flash file system:

```
switch: flash_init
```

Step 8 If you had set the console port speed to anything other than 9600, it has been reset to that particular speed. Change the emulation software line speed to match that of the switch console port.

- Step 9** Load any helper files:
switch: **load_helper**
- Step 10** Start the file transfer by using the XMODEM protocol.
switch: **copy xmodem: flash:image_filename.bin**
- Step 11** After the XMODEM request appears, use the appropriate command on the terminal-emulation software to start the transfer and to copy the software image into Flash memory.
- Step 12** Boot the newly-downloaded Cisco IOS image.
switch: **boot flash:image_filename.bin**
- Step 13** Use the **archive download-sw** privileged EXEC command to download the software image to the switch.
- Step 14** Use the **reload** privileged EXEC command to restart the switch and to verify that the new software image is operating properly.
- Step 15** Delete the `flash:image_filename.bin` file from the switch.
-

Recovering from a Lost or Forgotten Password

The default configuration for the switch allows an end user with physical access to the switch to recover from a lost password by interrupting the boot process during power-on and by entering a new password. These recovery procedures require that you have physical access to the switch.



Note

On these switches, a system administrator can disable some of the functionality of this feature by allowing an end user to reset a password only by agreeing to return to the default configuration. If you are an end user trying to reset a password when password recovery has been disabled, a status message shows this during the recovery process.

This section describes how to recover a forgotten or lost switch password. It provides two solutions:

- [Procedure with Password Recovery Enabled, page 34-4](#)
- [Procedure with Password Recovery Disabled, page 34-6](#)

You enable or disable password recovery by using the **service password-recovery** global configuration command.

Follow the steps in this procedure if you have forgotten or lost the switch password.

-
- Step 1** Connect a terminal or PC with terminal-emulation software to the switch console port.
- Step 2** Set the line speed on the emulation software to 9600 baud.
- Step 3** Power off the switch.

- Step 4** Press the **Mode** button, and at the same time, reconnect the power cord to the switch.
- You can release the **Mode** button a second or two after the LED above port 1 turns off. Several lines of information about the software appear with instructions, informing you if the password recovery procedure has been disabled or not.
- If you see a message that begins with this:
The system has been interrupted prior to initializing the flash file system. The following commands will initialize the flash file system

proceed to the [“Procedure with Password Recovery Enabled” section on page 34-4](#), and follow the steps.
 - If you see a message that begins with this:
The password-recovery mechanism has been triggered, but is currently disabled.

proceed to the [“Procedure with Password Recovery Disabled” section on page 34-6](#), and follow the steps.
- Step 5** After recovering the password, reload the switch:
- ```
Switch> reload
Proceed with reload? [confirm] y
```
- 

## Procedure with Password Recovery Enabled

If the password-recovery mechanism is enabled, this message appears:

The system has been interrupted prior to initializing the flash file system. The following commands will initialize the flash file system, and finish loading the operating system software:

```
flash_init
load_helper
boot
```

---

- Step 1** Initialize the Flash file system:
- ```
switch: flash_init
```
- Step 2** If you had set the console port speed to anything other than 9600, it has been reset to that particular speed. Change the emulation software line speed to match that of the switch console port.
- Step 3** Load any helper files:
- ```
switch: load_helper
```
- Step 4** Display the contents of Flash memory:
- ```
switch: dir flash:
```
- The switch file system appears in the directory.
- Step 5** Rename the configuration file to config.text.old.
- This file contains the password definition.
- ```
switch: rename flash:config.text flash:config.text.old
```



**Step 6** Boot the system:

```
switch: boot
```

You are prompted to start the setup program. Enter **N** at the prompt:

```
Continue with the configuration dialog? [yes/no]: N
```

**Step 7** At the switch prompt, enter privileged EXEC mode:

```
Switch> enable
```

**Step 8** Rename the configuration file to its original name:

```
Switch# rename flash:config.text.old flash:config.text
```

**Step 9** Copy the configuration file into memory:

```
Switch# copy flash:config.text system:running-config
Source filename [config.text]?
Destination filename [running-config]?
```

Press **Return** in response to the confirmation prompts.

The configuration file is now reloaded, and you can change the password.

**Step 10** Enter global configuration mode:

```
Switch# configure terminal
```

**Step 11** Change the password:

```
Switch (config)# enable secret password
```

The secret password can be from 1 to 25 alphanumeric characters, can start with a number, is case sensitive, and allows spaces but ignores leading spaces.

**Step 12** Return to privileged EXEC mode:

```
Switch (config)# exit
Switch#
```

**Step 13** Write the running configuration to the startup configuration file:

```
Switch# copy running-config startup-config
```

The new password is now in the startup configuration.

**Note**

This procedure is likely to leave your switch virtual interface in a shutdown state. You can see which interface is in this state by entering the **show running-config** privileged EXEC command. To re-enable the interface, enter the **interface vlan vlan-id** global configuration command, and specify the VLAN ID of the shutdown interface. With the switch in interface configuration mode, enter the **no shutdown** command.

**Step 14** Reload the switch:

```
Switch# reload
```

## Procedure with Password Recovery Disabled

If the password-recovery mechanism is disabled, this message appears:

The password-recovery mechanism has been triggered, but is currently disabled. Access to the boot loader prompt through the password-recovery mechanism is disallowed at this point. However, if you agree to let the system be reset back to the default system configuration, access to the boot loader prompt can still be allowed.

Would you like to reset the system back to the default configuration (y/n)?



### Caution

Returning the switch to the default configuration results in the loss of all existing configurations. We recommend that you contact your system administrator to verify if there are backup switch and VLAN configuration files.

- If you enter **n** (no), the normal boot process continues as if the **Mode** button had not been pressed; you cannot access the boot loader prompt, and you cannot enter a new password. You see the message:  
Press Enter to continue.....
- If you enter **y** (yes), the configuration file in Flash memory and the VLAN database file are deleted. When the default configuration loads, you can reset the password.

- 
- Step 1** Elect to continue with password recovery and lose the existing configuration:  
Would you like to reset the system back to the default configuration (y/n)? **Y**
- Step 2** Load any helper files:  
Switch: **load\_helper**
- Step 3** Display the contents of Flash memory:  
switch: **dir flash:**  
  
The switch file system appears in the directory.
- Step 4** Boot the system:  
Switch: **boot**  
  
You are prompted to start the setup program. To continue with password recovery, enter **N** at the prompt:  
Continue with the configuration dialog? [yes/no]: **N**
- Step 5** At the switch prompt, enter privileged EXEC mode:  
Switch> **enable**
- Step 6** Enter global configuration mode:  
Switch# **configure terminal**
- Step 7** Change the password:  
Switch (config)# **enable secret password**  
  
The secret password can be from 1 to 25 alphanumeric characters, can start with a number, is case sensitive, and allows spaces but ignores leading spaces.

**Step 8** Return to privileged EXEC mode:

```
Switch (config)# exit
Switch#
```

**Step 9** Write the running configuration to the startup configuration file:

```
Switch# copy running-config startup-config
```

The new password is now in the startup configuration.



**Note**

This procedure is likely to leave your switch virtual interface in a shutdown state. You can see which interface is in this state by entering the **show running-config** privileged EXEC command. To re-enable the interface, enter the **interface vlan** *vlan-id* global configuration command, and specify the VLAN ID of the shutdown interface. With the switch in interface configuration mode, enter the **no shutdown** command.

**Step 10** You must now reconfigure the switch. If the system administrator has the backup switch and VLAN configuration files available, you should use those.

## Preventing Autonegotiation Mismatches

The IEEE 802.3ab autonegotiation protocol manages the switch settings for speed (10 Mbps, 100 Mbps, and 1000 Mbps, excluding SFP module ports unless the 1000-BASE-T SFP is installed) 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.

# SFP Module Security and Identification

Cisco-approved small form-factor pluggable (SFP) modules have a serial EEPROM that contains the module serial number, the vendor name and ID, a unique security code, and cyclic redundancy check (CRC). When an SFP module is inserted in the switch, the switch software reads the EEPROM to verify the serial number, vendor name and vendor ID, and recompute the security code and CRC. If the serial number, the vendor name or vendor ID, the security code, or CRC is invalid, the software generates a security error message and places the interface in an error-disabled state.

**Note**

The security error message references the GBIC\_SECURITY facility. The switch supports SFP modules and does not support GBIC modules. Although the error message text refers to GBIC interfaces and modules, the security messages actually refer to the SFP modules and module interfaces. For more information about error messages, refer to the system message guide for this release.

If you are using a non-Cisco approved SFP module, remove the SFP module from the switch, and replace it with a Cisco-approved module. After inserting a Cisco-approved SFP module, use the **errdisable recovery cause gbic-invalid** global configuration command to verify the port status, and enter a time interval for recovering from the error-disabled state. After the elapsed interval, the switch brings the interface out of the error-disabled state and retries the operation. For more information about the **errdisable recovery** command, refer to the command reference for this release.

If the module is identified as a Cisco SFP module, but the system is unable to read vendor-data information to verify its accuracy, an SFP module error message is generated. In this case, you should remove and re-insert the SFP module. If it continues to fail, the SFP module might be defective.

## Using Ping

This section consists of this information:

- [Understanding Ping, page 34-8](#)
- [Executing Ping, page 34-9](#)

## 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 or have IP routing configured to route between those subnets. For more information, see [Chapter 28, “Configuring IP Unicast Routing.”](#)

IP routing is disabled by default on all switches. If you need to enable or configure IP routing, see [Chapter 28, “Configuring IP Unicast Routing.”](#)

Beginning in privileged EXEC mode, use this command to ping another device on the network from the switch:

| Command                              | Purpose                                                                         |
|--------------------------------------|---------------------------------------------------------------------------------|
| <b>ping ip</b> <i>host   address</i> | 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 34-1](#) describes the possible ping character output.

**Table 34-1 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, page 34-10](#)
- [Usage Guidelines, page 34-10](#)
- [Displaying the Physical Path, page 34-11](#)

## 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.

## 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 function properly, do not disable CDP. 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 19, “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 the 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 IP Traceroute

This section consists of this information:

- [Understanding IP Traceroute, page 34-11](#)
- [Executing IP Traceroute, page 34-12](#)

## Understanding IP Traceroute

You can use IP traceroute to identify the path that packets take through the network on a hop-by-hop basis. The command output displays all network layer (Layer 3) devices, such as routers, that the traffic passes through on the way to the destination.

Your switches can participate as the source or destination of the **traceroute** privileged EXEC command and might or might not appear as a hop in the **traceroute** command output. If the switch is the destination of the traceroute, it is displayed as the final destination in the traceroute output. Intermediate switches do not show up in the traceroute output if they are only bridging the packet from one port to another within the same VLAN. However, if the intermediate switch is a multilayer switch that is routing a particular packet, this switch shows up as a hop in the traceroute output.

The **traceroute** privileged EXEC command uses the Time To Live (TTL) field in the IP header to cause routers and servers to generate specific return messages. Traceroute starts by sending a User Datagram Protocol (UDP) datagram to the destination host with the TTL field set to 1. If a router finds a TTL value

of 1 or 0, it drops the datagram and sends back an Internet Control Message Protocol (ICMP) time-to-live-exceeded message to the sender. Traceroute determines the address of the first hop by examining the source address field of the ICMP time-to-live-exceeded message.

To identify the next hop, traceroute sends a UDP packet with a TTL value of 2. The first router decrements the TTL field by 1 and sends the datagram to the next router. The second router sees a TTL value of 1, discards the datagram, and returns the time-to-live-exceeded message to the source. This process continues until the TTL is incremented to a value large enough for the datagram to reach the destination host (or until the maximum TTL is reached).

To determine when a datagram reaches its destination, traceroute sets the UDP destination port number in the datagram to a very large value that the destination host is unlikely to be using. When a host receives a datagram destined to itself containing a destination port number that is unused locally, it sends an ICMP port unreachable error to the source. Because all errors except port unreachable errors come from intermediate hops, the receipt of a port unreachable error means this message was sent by the destination.

## Executing IP Traceroute

Beginning in privileged EXEC mode, follow this step to trace the path packets take through the network:

| Command                          | Purpose                                                      |
|----------------------------------|--------------------------------------------------------------|
| <b>traceroute ip</b> <i>host</i> | Trace the path packets take through the network by using IP. |



### Note

Though other protocol keywords are available with the **traceroute** privileged EXEC command, they are not supported in this release.

This example shows how to perform a **traceroute** to an IP host:

```
Switch# traceroute ip 171.9.15.10
```

```
Type escape sequence to abort.
```

```
Tracing the route to 171.69.115.10
```

```
 1 172.2.52.1 0 msec 0 msec 4 msec
 2 172.2.1.203 12 msec 8 msec 0 msec
 3 171.9.16.6 4 msec 0 msec 0 msec
 4 171.9.4.5 0 msec 4 msec 0 msec
 5 171.9.121.34 0 msec 4 msec 4 msec
 6 171.9.15.9 120 msec 132 msec 128 msec
 7 171.9.15.10 132 msec 128 msec 128 msec
```

```
Switch#
```

The display shows the hop count, IP address of the router, and the round-trip time in milliseconds for each of the three probes that are sent.

To terminate a trace in progress, 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 Debug Commands

This section explains how you use **debug** commands to diagnose and resolve internetworking problems. It contains this information:

- [Enabling Debugging on a Specific Feature, page 34-13](#)
- [Enabling All-System Diagnostics, page 34-14](#)
- [Redirecting Debug and Error Message Output, page 34-14](#)



## 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 Cisco technical support staff. 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 Switched Port Analyzer (SPAN):

```
Switch# debug span-session
```

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 you want to monitor. Use the **show running-config** command to check its configuration.
- Even if the switch is properly configured, it might not generate the type of traffic 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 span-session
```

Alternately, in privileged EXEC mode, you can enter the **undebug** form of the command:

```
Switch# undebug span-session
```

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 console. If you use this default, you can use a virtual terminal connection to monitor debug output instead of connecting to the console port.

Possible destinations include the console, 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 console 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 23, “Configuring System Message Logging.”](#)

## Using the show platform forward Command

The output from the **show platform forward** privileged EXEC command provides some useful information about the forwarding results if a packet entering an interface is sent through the system. Depending upon the parameters entered about the packet, the output provides lookup table results and port maps used to calculate forwarding destinations, bitmaps, and egress information.



### Note

For more syntax and usage information for the **show platform forward** command, refer to the switch command reference for this release.

Most of the information in the output from the command is useful mainly for technical support personnel, who have access to detailed information about the switch application-specific integrated circuits (ASICs). However, packet forwarding information can also be helpful in troubleshooting.

This is an example of the output from the **show platform forward** command on enhanced-services (ES) port 1 in VLAN 5 when the packet entering that port is addressed to unknown MAC addresses. The packet should be flooded to all other ports in VLAN 5.

```
Switch: show platform forward gigabitethernet1/1/1 vlan 10 1.1.1 2.2.2 ip 172.18.18.3
172.18.18.1 udp 10 20
```

```
Global Port Number:472, Asic Number:1
Src Real Vlan Id:10, Mapped Vlan Id:2
```

Ingress:

| Lookup                                            | Key-Used | Index-Hit | A-Data   |
|---------------------------------------------------|----------|-----------|----------|
| InptACL 40_AC121201_AC121203-00_40000014_000A0000 |          | 01FFA     | 03000000 |
| L2Local 80_00020002_00020002-00_00000000_00000000 |          | 01850     | 0000003A |

Station Descriptor:02D30000, DestIndex:02D5, RewriteIndex:F002

=====

Egress:Asic 0, switch 1

Output Packets:

-----

Packet 1

| Lookup                                             | Key-Used | Index-Hit | A-Data   |
|----------------------------------------------------|----------|-----------|----------|
| OutptACL 50_AC121201_AC121203-00_40000014_000A0000 |          | 01FFE     | 03000000 |

| Port    | Vlan | SrcMac         | DstMac         | Cos | Dscp |
|---------|------|----------------|----------------|-----|------|
| Gi1/0/1 | 0010 | 0001.0001.0001 | 0002.0002.0002 |     |      |

-----

Packet 2

| Lookup                                             | Key-Used | Index-Hit | A-Data   |
|----------------------------------------------------|----------|-----------|----------|
| OutptACL 50_AC121201_AC121203-00_40000014_000A0000 |          | 01FFE     | 03000000 |

| Port    | Vlan | SrcMac         | DstMac         | Cos | Dscp |
|---------|------|----------------|----------------|-----|------|
| Fa1/0/2 | 0010 | 0001.0001.0001 | 0002.0002.0002 |     |      |

-----

Packet 3

| Lookup                                             | Key-Used | Index-Hit | A-Data   |
|----------------------------------------------------|----------|-----------|----------|
| OutptACL 50_AC121201_AC121203-00_40000014_000A0000 |          | 01FFE     | 03000000 |

| Port    | Vlan | SrcMac         | DstMac         | Cos | Dscp |
|---------|------|----------------|----------------|-----|------|
| Fa1/0/3 | 0010 | 0001.0001.0001 | 0002.0002.0002 |     |      |

=====

Egress:Asic 1, switch 1

Output Packets:

-----

Packet 4

| Lookup                                             | Key-Used | Index-Hit | A-Data   |
|----------------------------------------------------|----------|-----------|----------|
| OutptACL 50_AC121201_AC121203-00_40000014_000A0000 |          | 01FFE     | 03000000 |

Packet dropped due to failed DEJA\_VU Check on Gi1/1/1

This is an example of the output when the packet coming in on ES port 1 in VLAN 5 is sent to an address already learned on the VLAN on another port. It should be forwarded from the port on which the address was learned.

```
Switch# show platform forward gigabitethernet1/1/1 vlan 5 1.1.1 0009.43a8.0145 ip 13.1.1.1
13.2.2.2 udp 10 20
```

```
Global Port Number:472, Asic Number:1
Src Real Vlan Id:5, Mapped Vlan Id:5
```

Ingress:

```

 Lookup Key-Used Index-Hit A-Data
InptACL 40_0D020202_0D010101-00_40000014_000A0000 01FFA 03000000
L2Local 80_00050009_43A80145-00_00000000_00000000 00086 02010197
Station Descriptor:F0050003, DestIndex:F005, RewriteIndex:0003

```

```
=====
```

```
Egress:Asic 3, switch 1
```

```
Output Packets:
```

```

```

```

Packet 1
 Lookup Key-Used Index-Hit A-Data
OutptACL 50_0D020202_0D010101-00_40000014_000A0000 01FFE 03000000

```

```

Port Vlan SrcMac DstMac Cos Dscpv
Fa1/0/5 0005 0001.0001.0001 0009.43A8.0145

```

This is an example of the output when the packet coming in on ES port 1 in VLAN 5 has a destination MAC address set to the router MAC address in VLAN 5 and the destination IP address unknown. Since there is no default route set, the packet should be dropped.

```
Switch# show platform forward gigabitethernet1/1/1 vlan 5 1.1.1 03.e319.ee44 ip 13.1.1.1 13.2.2.2 udp 10 20
```

```
Global Port Number:472, Asic Number:1
```

```
Src Real Vlan Id:5, Mapped Vlan Id:5
```

```
Ingress:
```

```

 Lookup Key-Used Index-Hit A-Data
InptACL 40_0D020202_0D010101-00_41000014_000A0000 01FFA 03000000
L3Local 00_00000000_00000000-90_00001400_0D020202 010F0 01880290
L3Scndr 12_0D020202_0D010101-00_40000014_000A0000 034E0 000C001D_00000000
Lookup Used:Secondary
Station Descriptor:02260000, DestIndex:0226, RewriteIndex:0000

```

This is an example of the output when the packet coming in on ES port 1 in VLAN 5 has a destination MAC address set to the router MAC address in VLAN 5 and the destination IP address set to an IP address that is in the IP routing table. It should be forwarded as specified in the routing table.

```
Switch# show platform forward gigabitethernet1/1/1 vlan 5 1.1.1 03.e319.ee44 ip 110.1.5.5 16.1.10.5
```

```
Global Port Number:472, Asic Number:1
```

```
Src Real Vlan Id:5, Mapped Vlan Id:5
```

```
Ingress:
```

```

 Lookup Key-Used Index-Hit A-Data
InptACL 40_10010A05_0A010505-00_41000014_000A0000 01FFA 03000000
L3Local 00_00000000_00000000-90_00001400_10010A05 010F0 01880290
L3Scndr 12_10010A05_0A010505-00_40000014_000A0000 01D28 30090001_00000000
Lookup Used:Secondary
Station Descriptor:F0070007, DestIndex:F007, RewriteIndex:0007

```

```
=====
```

```
Egress:Asic 3, switch 1e
```

```
Output Packets:
```

```

```

```

Packet 1
 Lookup Key-Used Index-Hit A-Data
OutptACL 50_10010A05_0A010505-00_40000014_000A0000 01FFE 03000000

```

```

Port Vlan SrcMac DstMac Cos Dscpv
Gi1/0/1 0007 XXXX.XXXX.0246 0009.43A8.0147

```

## Using the crashinfo File

The crashinfo file saves information that helps Cisco technical support representatives to debug problems that caused the Cisco IOS image to fail (crash). The switch writes the crash information to the console at the time of the failure, and the file is created the next time you boot the Cisco IOS image after the failure (instead of while the system is failing).

The information in the file includes the Cisco IOS image name and version that failed, a list of the processor registers, and a stack trace. You can provide this information to the Cisco 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.





## Supported MIBs

---

This appendix lists the supported MIBs for this release on the Catalyst 3750 Metro switch. It contains these sections:

- [MIB List, page A-1](#)
- [Using FTP to Access the MIB Files, page A-3](#)

## MIB List

- BRIDGE-MIB (RFC1493)

**Note**

The BRIDGE-MIB supports the context of a single VLAN. By default, SNMP messages using the configured community string always provide information for VLAN 1. To obtain the BRIDGE-MIB information for other VLANs, for example VLAN x, use this community string in the SNMP message: configured community string @x.

- CISCO-CDP-MIB
- CISCO-CLASS-BASED-QOS-MIB
- CISCO-CLUSTER-MIB
- CISCO-CONFIG-COPY-MIB
- CISCO-CONFIG-MAN-MIB
- CISCO-ENTITY-FRU-CONTROL-MIB
- CISCO-ENVMON-MIB
- CISCO-FLASH-MIB (Flash memory on all switches is modeled as removable Flash memory.)
- CISCO-FTP-CLIENT-MIB
- CISCO-HSRP-MIB
- CISCO-HSRP-EXT-MIB (partial support)
- CISCO-IGMP-FILTER-MIB
- CISCO-IMAGE-MIB
- CISCO IP-STAT-MIB
- CISCO-L2L3-INTERFACE-MIB

- CISCO-LACP-MIB
- CISCO-MAC-NOTIFICATION-MIB
- CISCO-MEMORY-POOL-MIB
- CISCO-PAGP-MIB
- CISCO-PING-MIB
- CISCO-PROCESS-MIB
- CISCO-RTTMON-MIB
- CISCO-STACKMAKER-MIB
- CISCO-STP-EXTENSIONS-MIB
- CISCO-SYSLOG-MIB
- CISCO-TCP-MIB
- CISCO-VLAN-IFTABLE-RELATIONSHIP-MIB
- CISCO-VLAN-MEMBERSHIP-MIB
- CISCO-VTP-MIB
- ENTITY-MIB
- ETHERLIKE\_MIB
- IEEE8023-LACP-MIB
- IF-MIB (In and out counters for VLANs are not supported.)
- IGMP-MIB
- IPMROUTE-MIB
- MPLS-LDP-MIB
- MPLS-LSR-MIB
- MPLS-VPN-MIB
- OLD-CISCO-CHASSIS-MIB
- OLD-CISCO-FLASH-MIB
- OLD-CISCO-INTERFACES-MIB
- OLD-CISCO-IP-MIB
- OLD-CISCO-SYS-MIB
- OLD-CISCO-TCP-MIB
- OLD-CISCO-TS-MIB
- PIM-MIB
- RFC1213-MIB (Functionality is as per the agent capabilities specified in the CISCO-RFC1213-CAPABILITY.my.)
- RFC1253-MIB (OSPF-MIB)
- RMON-MIB
- RMON2-MIB
- SNMP-FRAMEWORK-MIB
- SNMP-MPD-MIB



- SNMP-NOTIFICATION-MIB
- SNMP-TARGET-MIB
- SNMPv2-MIB
- TCP-MIB
- UDP-MIB

**Note**

You can also use this URL for a list of supported MIBs for the Catalyst 3750 Metro switch:  
<ftp://ftp.cisco.com/pub/mibs/supportlists/cat3750me/cat3750me-supportlist.html>  
You can access other information about MIBs and Cisco products on the Cisco web site:  
<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>

## Using FTP to Access the MIB Files

You can obtain each MIB file by using this procedure:

- 
- |               |                                                                                                   |
|---------------|---------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Use FTP to access the server <b>ftp.cisco.com</b> .                                               |
| <b>Step 2</b> | Log in with the username <b>anonymous</b> .                                                       |
| <b>Step 3</b> | Enter your e-mail username when prompted for the password.                                        |
| <b>Step 4</b> | At the <b>ftp&gt;</b> prompt, change directories to <b>/pub/mibs/v1</b> and <b>/pub/mibs/v2</b> . |
| <b>Step 5</b> | Use the <b>get MIB_filename</b> command to obtain a copy of the MIB file.                         |
-





## Working with the Cisco IOS File System, Configuration Files, and Software Images

This appendix describes how to manipulate the Catalyst 3750 Metro switch flash file system, how to copy configuration files, and how to archive (upload and download) software images to a 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 IOS Configuration Fundamentals Command Reference for Release 12.1*.

This appendix consists of these sections:

- [Working with the Flash File System, page B-1](#)
- [Working with Configuration Files, page B-8](#)
- [Working with Software Images, page B-19](#)

## Working with the Flash File System

The flash file system is a single flash device on which you can store files. It also provides several commands to help you manage software image and configuration files. The default flash file system on the switch is named *flash:*.

This section contains this information:

- [Displaying Available File Systems, page B-2](#)
- [Setting the Default File System, page B-3](#)
- [Displaying Information about Files on a File System, page B-3](#)
- [Creating and Removing Directories, page B-4](#)
- [Copying Files, page B-4](#)
- [Deleting Files, page B-5](#)
- [Creating, Displaying, and Extracting tar Files, page B-5](#)
- [Displaying the Contents of a File, page B-7](#)

## 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
* 15998976 5135872 flash rw flash:
 - - opaque rw bs:
 - - opaque rw vb:
 524288 520138 nvram rw nvram:
 - - network rw tftp:
 - - opaque rw null:
 - - opaque rw system:
 - - opaque ro xmodem:
 - - opaque ro ymodem:
```

**Table B-1** 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.<br><b>flash</b> —The file system is for a flash memory device.<br><b>nvram</b> —The file system is for a NVRAM device.<br><b>opaque</b> —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.<br><b>unknown</b> —The file system is an unknown type.                                                                                                                                                                                    |
| Flags    | Permission for file system.<br><b>ro</b> —read-only.<br><b>rw</b> —read/write.<br><b>wo</b> —write-only.                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Prefixes | Alias for file system.<br><b>flash:</b> —Flash file system.<br><b>nvram:</b> —NVRAM.<br><b>null:</b> —Null destination for copies. You can copy a remote file to null to find its size.<br><b>rcp:</b> —Remote copy (RCP) network server.<br><b>system:</b> —Contains the system memory, including the running configuration.<br><b>tftp:</b> —TFTP network server.<br><b>xmodem:</b> —Obtain the file from a network machine by using the Xmodem protocol.<br><b>ymodem:</b> —Obtain the file from a network machine by using the Ymodem protocol. |

## 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 B-2](#):

**Table B-2** Commands for Displaying Information About Files

| Command                                   | Description                                                                                                                                                                |
|-------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>dir</b> [/all] [filesystem:][filename] | Display a list of files on a file system.                                                                                                                                  |
| <b>show file systems</b>                  | Display more information about each of the files on a file system.                                                                                                         |
| <b>show file information</b> file-url     | Display information about a specific file.                                                                                                                                 |
| <b>show file descriptors</b>              | 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.

|        | Command                | Purpose                                                                                                                                |
|--------|------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | <b>dir filesystem:</b> | Display the directories on the specified file system.<br>For <i>filesystem:</i> , use <b>flash:</b> for the system board flash device. |
| Step 2 | <b>cd new_configs</b>  | Change to the directory of interest.<br>The command example shows how to change to the directory named <i>new_configs</i> .            |
| Step 3 | <b>pwd</b>             | Display the working directory.                                                                                                         |

## Creating and Removing Directories

Beginning in privileged EXEC mode, follow these steps to create and remove a directory:

|        | Command                         | Purpose                                                                                                                                                                                                                                                                                                                          |
|--------|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | <b>dir</b> <i>filesystem:</i>   | Display the directories on the specified file system.<br>For <i>filesystem:</i> , use <b>flash:</b> for the system board flash device.                                                                                                                                                                                           |
| Step 2 | <b>mkdir</b> <i>old_configs</i> | Create a new directory.<br>The command example shows how to create the directory named <i>old_configs</i> .<br>Directory names are case sensitive.<br>Directory names are limited to 45 characters between the slashes (/); the name cannot contain control characters, spaces, deletes, slashes, quotes, semicolons, or colons. |
| Step 3 | <b>dir</b> <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 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 from special file systems (**xmodem:**, **ymodem:**) as the source 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:

- FTP—**ftp:**[[//username [:password]@location]/directory]/filename
- RCP—**rnp:**[[//username@location]/directory]/filename
- 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 B-8](#).

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 B-19](#).

## 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.

When you attempt to delete any files, 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.



### Note

Instead of using the **copy** privileged EXEC command or the **archive tar** privileged EXEC command, we recommend using the **archive download-sw** and **archive upload-sw** privileged EXEC commands to download and upload software image files.

## Creating a tar File

To create a tar file and write files into it, use this 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 FTP, the syntax is **ftp:[[/username[:password]@location]/directory]/tar-filename.tar**
- For the RCP, the syntax is **rcp:[[/username@location]/directory]/tar-filename.tar**
- For the 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 FTP, the syntax is **ftp:[[/username[:password]@location]/directory]/tar-filename.tar**
- For the RCP, the syntax is **rcp:[[/username@location]/directory]/tar-filename.tar**
- For the 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 those files appear. If none are specified, all files and directories appear.

This example shows how to display the contents of a switch tar file that is in flash memory:

```
Switch# archive tar /table flash:image-tv0-m.tar
info (219 bytes)
image-tv0-mz-121/ (directory)
image-tv0-mz-121/html/ (directory)
image-tv0-mz-121/html/foo.html (0 bytes)
image-tv0-mz-121/image-tv0-mz-121.bin (610856 bytes)
image-tv0-mz-121/info (219 bytes)
info.ver (219 bytes)
```



This example shows how to display only the *image-tv0-mz-121/html* directory and its contents:

```
Switch# archive tar /table flash:image-tv0-m.tar image-tv0-mz-121/html
image-tv0-mz-121/html/ (directory)
image-tv0-mz-121/html/foo.html (0 bytes)
```

## Extracting a tar File

To extract a tar file into a directory on the flash file system, use this privileged EXEC command:

```
archive tar /xtract source-url flash:/file-url [dir/file...]
```

For *source-url*, specify the source URL alias for the local file system. These options are supported:

- For the local flash file system, the syntax is **flash:**
- For the FTP, the syntax is **ftp:[[/username[:password]@location]/directory]/tar-filename.tar**
- For the RCP, the syntax is **rcp:[[/username@location]/directory]/tar-filename.tar**
- For the 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 [dir/file...]**, specify the location on the local flash file system into which the tar file is extracted. Use the *dir/file...* option to specify an optional list of files or directories within the tar file to be extracted. 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 linenummer
service udp-small-servers
service pt-vty-logging
!

<output truncated>
```

# Working with Configuration Files

This section describes how to create, load, and maintain configuration files.

Configuration files contain commands entered to customize the function of the Cisco IOS software. A way to create a basic configuration file is to use the **setup** program or to enter the **setup** privileged EXEC command. For more information, see [Chapter 3, “Assigning the Switch IP Address and Default Gateway.”](#)

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 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 TCP/IP stack, which is connection-oriented.

This section includes this information:

- [Guidelines for Creating and Using Configuration Files, page B-8](#)
- [Configuration File Types and Location, page B-9](#)
- [Creating a Configuration File By Using a Text Editor, page B-9](#)
- [Copying Configuration Files By Using TFTP, page B-10](#)
- [Copying Configuration Files By Using FTP, page B-12](#)
- [Copying Configuration Files By Using RCP, page B-15](#)
- [Clearing Configuration Information, page B-18](#)

## 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 console port for the initial configuration of the switch. If you are accessing the switch through a network connection instead of through a direct connection to the console port, keep in mind that some configuration changes (such as changing the switch IP address or disabling ports) can cause a loss of connectivity to the switch.
- If no password has been set on the switch, we recommend that you set one by using the **enable secret** *secret-password* global configuration command.

**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:

- 
- Step 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 B-11, the [“Downloading a Configuration File By Using FTP”](#) section on page B-13, or the [“Downloading a Configuration File By Using RCP”](#) section on page B-16.
- Step 2** Open the configuration file in a text editor, such as vi or emacs on UNIX or Notepad on a PC.
- Step 3** Extract the portion of the configuration file with the desired commands, and save it in a new file.

- Step 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).
- Step 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, page B-10](#)
- [Downloading the Configuration File By Using TFTP, page B-11](#)
- [Uploading the Configuration File By Using TFTP, page B-11](#)

## 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:

- 
- Step 1** Copy the configuration file to the appropriate TFTP directory on the workstation.
  - Step 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 B-10.
  - Step 3** Log into the switch through the console port or a Telnet session.
  - Step 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-confg* 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:

- 
- Step 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 B-10.
  - Step 2** Log into the switch through the console port or a Telnet session.
  - Step 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 resides 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, page B-12](#)
- [Downloading a Configuration File By Using FTP, page B-13](#)
- [Uploading a Configuration File By Using FTP, page B-14](#)

## 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 the console or 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:

|        | Command                                                                                                                                                                                                                                               | Purpose                                                                                                                                                                         |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 |                                                                                                                                                                                                                                                       | Verify that the FTP server is properly configured by referring to the <a href="#">“Preparing to Download or Upload a Configuration File By Using FTP”</a> section on page B-12. |
| Step 2 |                                                                                                                                                                                                                                                       | Log into the switch through the console port or a Telnet session.                                                                                                               |
| Step 3 | <b>configure terminal</b>                                                                                                                                                                                                                             | Enter global configuration mode on the switch.<br><br>This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).               |
| Step 4 | <b>ip ftp username</b> <i>username</i>                                                                                                                                                                                                                | (Optional) Change the default remote username.                                                                                                                                  |
| Step 5 | <b>ip ftp password</b> <i>password</i>                                                                                                                                                                                                                | (Optional) Change the default password.                                                                                                                                         |
| Step 6 | <b>end</b>                                                                                                                                                                                                                                            | Return to privileged EXEC mode.                                                                                                                                                 |
| Step 7 | <b>copy</b><br><b>ftp:[[/[username[:password]@]location]/directory]</b><br><b>/filename] system:running-config</b><br><br>or<br><br><b>copy</b><br><b>ftp:[[/[username[:password]@]location]/directory]</b><br><b>/filename] nvram:startup-config</b> | 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-conf* 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-conf system:running-config
Configure using host1-conf from 172.16.101.101? [confirm]
Connected to 172.16.101.101
Loading 1112 byte file host1-conf:![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: nvram: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:

|        | Command                                                                                                                                                                                                                                                 | Purpose                                                                                                                                                                         |
|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 |                                                                                                                                                                                                                                                         | Verify that the FTP server is properly configured by referring to the <a href="#">“Preparing to Download or Upload a Configuration File By Using FTP”</a> section on page B-12. |
| Step 2 |                                                                                                                                                                                                                                                         | Log into the switch through the console port or a Telnet session.                                                                                                               |
| Step 3 | <b>configure terminal</b>                                                                                                                                                                                                                               | Enter global configuration mode.<br><br>This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).                             |
| Step 4 | <b>ip ftp username</b> <i>username</i>                                                                                                                                                                                                                  | (Optional) Change the default remote username.                                                                                                                                  |
| Step 5 | <b>ip ftp password</b> <i>password</i>                                                                                                                                                                                                                  | (Optional) Change the default password.                                                                                                                                         |
| Step 6 | <b>end</b>                                                                                                                                                                                                                                              | Return to privileged EXEC mode.                                                                                                                                                 |
| Step 7 | <b>copy system:running-config</b><br><b>ftp:[[[/[username[:password]@]location]/directory]</b><br><b>/filename]</b><br><br>or<br><br><b>copy nvram:startup-config</b><br><b>ftp:[[[/[username[:password]@]location]/directory]</b><br><b>/filename]</b> | 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 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, page B-16](#)
- [Downloading a Configuration File By Using RCP, page B-16](#)
- [Uploading a Configuration File By Using RCP, page B-17](#)

## 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 the console or 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:

|        | Command                                 | Purpose                                                                                                                                                                         |
|--------|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 |                                         | Verify that the RCP server is properly configured by referring to the <a href="#">“Preparing to Download or Upload a Configuration File By Using RCP”</a> section on page B-16. |
| Step 2 |                                         | Log into the switch through the console port or a Telnet session.                                                                                                               |
| Step 3 | <b>configure terminal</b>               | Enter global configuration mode.<br><br>This step is required only if you override the default remote username (see Steps 4 and 5).                                             |
| Step 4 | <b>ip rcmd remote-username username</b> | (Optional) Specify the remote username.                                                                                                                                         |

|        | Command                                                                                                                                                                                                                      | Purpose                                                                                                                         |
|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| Step 5 | <b>end</b>                                                                                                                                                                                                                   | Return to privileged EXEC mode.                                                                                                 |
| Step 6 | <b>copy</b><br><b>rcp:[[/[username@]location]/directory]/filename]</b><br><b>system:running-config</b><br><br>or<br><br><b>copy</b><br><b>rcp:[[/[username@]location]/directory]/filename]</b><br><b>nvrn:startup-config</b> | 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-config* 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: nvrn: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 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:

|        | Command                   | Purpose                                                                                                                                                                         |
|--------|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 |                           | Verify that the RCP server is properly configured by referring to the <a href="#">“Preparing to Download or Upload a Configuration File By Using RCP”</a> section on page B-16. |
| Step 2 |                           | Log into the switch through the console port or a Telnet session.                                                                                                               |
| Step 3 | <b>configure terminal</b> | Enter global configuration mode.<br><br>This step is required only if you override the default remote username (see Steps 4 and 5).                                             |

|        | Command                                                                                                                                                                                           | Purpose                                                                                                         |
|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| Step 4 | <b>ip rcmd remote-username</b> <i>username</i>                                                                                                                                                    | (Optional) Specify the remote username.                                                                         |
| Step 5 | <b>end</b>                                                                                                                                                                                        | Return to privileged EXEC mode.                                                                                 |
| Step 6 | <b>copy system:running-config</b><br><b>rcp:[[/[username@]location]/directory]/filename]</b><br>or<br><b>copy nvram:startup-config</b><br><b>rcp:[[/[username@]location]/directory]/filename]</b> | 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-config* 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-config
Write file switch-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:

```
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. If you reboot the switch with no startup configuration, the switch enters the setup program so that you can reconfigure the switch with all new settings.

### 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 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 and the Cisco IOS code.

**Note**

Instead of using the **copy** privileged EXEC command or the **archive tar** privileged EXEC command, we recommend using the **archive download-sw** and **archive upload-sw** privileged EXEC commands to download and upload software image 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 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 TCP/IP stack, which is connection-oriented.

This section includes this information:

- [Image Location on the Switch, page B-20](#)
- [tar File Format of Images on a Server or Cisco.com, page B-20](#)
- [Copying Image Files By Using TFTP, page B-21](#)
- [Copying Image Files By Using FTP, page B-24](#)
- [Copying Image Files By Using RCP, page B-28](#)

**Note**

For a list of software images and the supported upgrade paths, refer to the release notes that shipped with your switch.

## Image Location on the Switch

The Cisco IOS image is stored as a `.bin` file in a directory that shows the version number. 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:

- An *info* file, which serves as a table of contents for the tar file
- One or more subdirectories containing other images and files, such as Cisco IOS images

This example shows some of the information contained in the info file. [Table B-3](#) provides additional details about this information:

```
version_suffix:i5-121.14-AX
version_directory:c3750me-i5-mz.121.14-AX
image_system_type_id:0x00000001
image_name:c3750me-i5-mz.121.14-AX.bin
ios_image_file_size:4782592
total_image_file_size:4784128
image_feature:LAYER_3|MIN_DRAM_MEG=64
image_family:C3750ME
stacking_number:1.1
board_ids:0x00000004
info_end:
```



### Note

Disregard the `stacking_number` field. It does not apply to the switch.

**Table B-3** info File Description

| Field                 | Description                                                                                                                                                         |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| version_suffix        | Specifies the Cisco IOS image version string suffix                                                                                                                 |
| version_directory     | Specifies the directory where the Cisco IOS image is 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 memory is required to hold just the Cisco IOS image           |
| total_image_file_size | Specifies the size of all the images (the Cisco IOS image files) in the tar file, which is an approximate measure of how much flash memory 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                                                                                             |

## 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.

**Note**

Instead of using the **copy** privileged EXEC command or the **archive tar** privileged EXEC command, we recommend using the **archive download-sw** and **archive upload-sw** privileged EXEC commands to download and upload software image files.

This section includes this information:

- [Preparing to Download or Upload an Image File By Using TFTP, page B-21](#)
- [Downloading an Image File By Using TFTP, page B-22](#)
- [Uploading an Image File By Using TFTP, page B-23](#)

## 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.

|        | Command                                                                                       | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------|-----------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 |                                                                                               | Copy the image to the appropriate TFTP directory on the workstation. Make sure the TFTP server is properly configured; see the <a href="#">“Preparing to Download or Upload an Image File By Using TFTP”</a> section on page B-21.                                                                                                                                                                                                                                                                                                                                                                                                               |
| Step 2 |                                                                                               | Log into the switch through the console port or a Telnet session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Step 3 | <b>archive download-sw /overwrite /reload tftp://[//location]/directory/image-name.tar</b>    | Download the image file from the TFTP server to the switch, and overwrite the current image. <ul style="list-style-type: none"> <li>• The <b>/overwrite</b> option overwrites the software image in flash memory with the downloaded image.</li> <li>• The <b>/reload</b> option reloads the system after downloading the image unless the configuration has been changed and not been saved.</li> <li>• For <i>//location</i>, specify the IP address of the TFTP server.</li> <li>• For <i>/directory/image-name.tar</i>, specify the directory (optional) and the image to download. Directory and image names are case sensitive.</li> </ul> |
| Step 4 | <b>archive download-sw /leave-old-sw /reload tftp://[//location]/directory/image-name.tar</b> | Download the image file from the TFTP server to the switch, and keep the current image. <ul style="list-style-type: none"> <li>• The <b>/leave-old-sw</b> option keeps the old software version after a download.</li> <li>• The <b>/reload</b> option reloads the system after downloading the image unless the configuration has been changed and not been saved.</li> <li>• For <i>//location</i>, specify the IP address of the TFTP server.</li> <li>• For <i>/directory/image-name.tar</i>, specify the directory (optional) and the image to download. Directory and image names are case sensitive.</li> </ul>                           |

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it aborts 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.

Beginning in privileged EXEC mode, follow these steps to upload an image to a TFTP server:

|        | Command                                                                        | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------|--------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 |                                                                                | Make sure the TFTP server is properly configured; see the <a href="#">“Preparing to Download or Upload an Image File By Using TFTP” section on page B-21</a> .                                                                                                                                                                                                                                                                                             |
| Step 1 |                                                                                | Log into the switch through the console port or a Telnet session.                                                                                                                                                                                                                                                                                                                                                                                          |
| Step 2 | <b>archive upload-sw</b><br><b>tftp:[[/location]/directory]/image-name.tar</b> | Upload the currently running switch image to the TFTP server. <ul style="list-style-type: none"> <li>For <i>//location</i>, specify the IP address of the TFTP server.</li> <li>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.</li> </ul> |

The **archive upload-sw** privileged EXEC command builds an image file on the server by uploading the the info file and the Cisco IOS image file. 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.



### Note

Instead of using the **copy** privileged EXEC command or the archive tar privileged EXEC command, we recommend using the **archive download-sw** and **archive upload-sw** privileged EXEC commands to download and upload software image files.

This section includes this information:

- [Preparing to Download or Upload an Image File By Using FTP, page B-24](#)
- [Downloading an Image File By Using FTP, page B-25](#)
- [Uploading an Image File By Using FTP, page B-27](#)

## 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 the console or 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.

|        | Command                                | Purpose                                                                                                                                                                  |
|--------|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 |                                        | Verify that the FTP server is properly configured by referring to the <a href="#">“Preparing to Download or Upload an Image File By Using FTP”</a> section on page B-24. |
| Step 2 |                                        | Log into the switch through the console port or a Telnet session.                                                                                                        |
| Step 3 | <b>configure terminal</b>              | Enter global configuration mode.<br><br>This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).                      |
| Step 4 | <b>ip ftp username</b> <i>username</i> | (Optional) Change the default remote username.                                                                                                                           |
| Step 5 | <b>ip ftp password</b> <i>password</i> | (Optional) Change the default password.                                                                                                                                  |
| Step 6 | <b>end</b>                             | Return to privileged EXEC mode.                                                                                                                                          |

| Command                                                                                                                                               | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 7</b><br><b>archive download-sw /overwrite /reload</b><br><b>ftp:[[/username[:password]@location]/directory]</b><br><b>/image-name.tar</b>    | Download the image file from the FTP server to the switch, and overwrite the current image. <ul style="list-style-type: none"> <li>• The <b>/overwrite</b> option overwrites the software image in flash memory with the downloaded image.</li> <li>• The <b>/reload</b> option reloads the system after downloading the image unless the configuration has been changed and not been saved.</li> <li>• For <b>//username[:password]</b>, specify the username and password; these must be associated with an account on the FTP server. For more information, see the <a href="#">“Preparing to Download or Upload an Image File By Using FTP”</a> section on page B-24.</li> <li>• For <b>@location</b>, specify the IP address of the FTP server.</li> <li>• For <b>directory/image-name.tar</b>, specify the directory (optional) and the image to download. Directory and image names are case sensitive.</li> </ul> |
| <b>Step 8</b><br><b>archive download-sw /leave-old-sw /reload</b><br><b>ftp:[[/username[:password]@location]/directory]</b><br><b>/image-name.tar</b> | Download the image file from the FTP server to the switch, and keep the current image. <ul style="list-style-type: none"> <li>• The <b>/leave-old-sw</b> option keeps the old software version after a download.</li> <li>• The <b>/reload</b> option reloads the system after downloading the image unless the configuration has been changed and not been saved.</li> <li>• For <b>//username[:password]</b>, specify the username and password. These must be associated with an account on the FTP server. For more information, see the <a href="#">“Preparing to Download or Upload an Image File By Using FTP”</a> section on page B-24.</li> <li>• For <b>@location</b>, specify the IP address of the FTP server.</li> <li>• For <b>directory/image-name.tar</b>, specify the directory (optional) and the image to download. Directory and image names are case sensitive.</li> </ul>                           |

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it aborts 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.

Beginning in privileged EXEC mode, follow these steps to upload an image to an FTP server:

|        | Command                                                                                                        | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------|----------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 |                                                                                                                | Verify that the FTP server is properly configured by referring to the <a href="#">“Preparing to Download or Upload a Configuration File By Using FTP”</a> section on page B-12.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Step 2 |                                                                                                                | Log into the switch through the console port or a Telnet session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Step 3 | <b>configure terminal</b>                                                                                      | Enter global configuration mode.<br><br>This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Step 4 | <b>ip ftp username</b> <i>username</i>                                                                         | (Optional) Change the default remote username.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Step 5 | <b>ip ftp password</b> <i>password</i>                                                                         | (Optional) Change the default password.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Step 6 | <b>end</b>                                                                                                     | Return to privileged EXEC mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Step 7 | <b>archive upload-sw</b><br><b>ftp:[[/[username[:password]@]location]/directory]/</b><br><b>image-name.tar</b> | Upload the currently running switch image to the FTP server. <ul style="list-style-type: none"> <li>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 <a href="#">“Preparing to Download or Upload an Image File By Using FTP”</a> section on page B-24.</li> <li>For <i>@location</i>, specify the IP address of the FTP server.</li> <li>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.</li> </ul> |

The **archive upload-sw** command builds an image file on the server by uploading the info file and the Cisco IOS image file. 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.

**Note**

Instead of using the **copy** privileged EXEC command or the **archive tar** privileged EXEC command, we recommend using the **archive download-sw** and **archive upload-sw** privileged EXEC commands to download and upload software image files.

This section includes this information:

- [Preparing to Download or Upload an Image File By Using RCP, page B-28](#)
- [Downloading an Image File By Using RCP, page B-29](#)
- [Uploading an Image File By Using RCP, page B-31](#)

## 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 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 the console or 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.

|        | Command | Purpose                                                                                                                                                                  |
|--------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 |         | Verify that the RCP server is properly configured by referring to the <a href="#">“Preparing to Download or Upload an Image File By Using RCP”</a> section on page B-28. |
| Step 2 |         | Log into the switch through the console port or a Telnet session.                                                                                                        |

|        | Command                                                                                                            | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------|--------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 3 | <b>configure terminal</b>                                                                                          | Enter global configuration mode.<br><br>This step is required only if you override the default remote username (see Steps 4 and 5).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Step 4 | <b>ip rcmd remote-username</b> <i>username</i>                                                                     | (Optional) Specify the remote username.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Step 5 | <b>end</b>                                                                                                         | Return to privileged EXEC mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Step 6 | <b>archive download-sw /overwrite /reload</b><br><b>rcp:[[//[username@]location]/directory]/image-name.tar]</b>    | Download the image file from the RCP server to the switch, and overwrite the current image. <ul style="list-style-type: none"> <li>• The <b>/overwrite</b> option overwrites the software image in flash memory with the downloaded image.</li> <li>• The <b>/reload</b> option reloads the system after downloading the image unless the configuration has been changed and not been saved.</li> <li>• For <b>//username</b>, 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 <a href="#">“Preparing to Download or Upload an Image File By Using RCP”</a> section on page B-28.</li> <li>• For <b>@location</b>, specify the IP address of the RCP server.</li> <li>• For <b>/directory/image-name.tar</b>, specify the directory (optional) and the image to download. Directory and image names are case sensitive.</li> </ul> |
| Step 7 | <b>archive download-sw /leave-old-sw /reload</b><br><b>rcp:[[//[username@]location]/directory]/image-name.tar]</b> | Download the image file from the RCP server to the switch, and keep the current image. <ul style="list-style-type: none"> <li>• The <b>/leave-old-sw</b> option keeps the old software version after a download.</li> <li>• The <b>/reload</b> option reloads the system after downloading the image unless the configuration has been changed and not been saved.</li> <li>• For <b>//username</b>, 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 <a href="#">“Preparing to Download or Upload an Image File By Using RCP”</a> section on page B-28.</li> <li>• For <b>@location</b>, specify the IP address of the RCP server.</li> <li>• For <b>/directory/image-name.tar</b>, specify the directory (optional) and the image to download. Directory and image names are case sensitive.</li> </ul>                                        |

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it aborts 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 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 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.

Beginning in privileged EXEC mode, follow these steps to upload an image to an RCP server:

|        | Command                                                                                    | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------|--------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 |                                                                                            | Verify that the RCP server is properly configured by referring to the <a href="#">“Preparing to Download or Upload an Image File By Using RCP”</a> section on page B-28.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Step 2 |                                                                                            | Log into the switch through the console port or a Telnet session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Step 3 | <b>configure terminal</b>                                                                  | Enter global configuration mode.<br><br>This step is required only if you override the default remote username (see Steps 4 and 5).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Step 4 | <b>ip rcmd remote-username</b> <i>username</i>                                             | (Optional) Specify the remote username.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Step 5 | <b>end</b>                                                                                 | Return to privileged EXEC mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Step 6 | <b>archive upload-sw</b><br><b>rtp: [[[/username@]location]/directory]/image-name.tar]</b> | Upload the currently running switch image to the RCP server. <ul style="list-style-type: none"> <li>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 <a href="#">“Preparing to Download or Upload an Image File By Using RCP”</a> section on page B-28.</li> <li>For <i>@location</i>, specify the IP address of the RCP server.</li> <li>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.</li> <li>The <i>image-name.tar</i> is the name of software image to be stored on the server.</li> </ul> |

The **archive upload-sw** privileged EXEC command builds an image file on the server by uploading the info file and the Cisco IOS image file. 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.

---



## Unsupported Commands in Cisco IOS Release 12.1(14)AX

---

This appendix lists some of the command-line interface (CLI) commands that are displayed when you enter the question mark (?) at the Catalyst 3750 Metro switch prompt but are not supported in this release, either because they are not tested, or because of Catalyst 3750 Metro switch hardware limitations. This is not a complete list. The unsupported commands are listed by software feature and command mode.

### Access Control Lists

#### Unsupported Privileged EXEC Commands

**access-enable** [*host*] [*timeout minutes*]

**access-template** [*access-list-number* | *name*] [*dynamic-name*] [*source*] [*destination*] [**timeout** *minutes*]

**clear access-template** [*access-list-number* | *name*] [*dynamic-name*] [*source*] [*destination*].

#### Unsupported Global Configuration Commands

**access-list rate-limit** *acl-index* {*precedence* | **mask** *prec-mask*}

**access-list dynamic extended**

### ARP

#### Unsupported Global Configuration Commands

**arp** *ip-address hardware-address smds*

**arp** *ip-address hardware-address srp-a*

**arp** *ip-address hardware-address srp-b*

## Unsupported Interface Configuration Commands

**arp probe**  
**ip probe proxy**

## Debug

### Unsupported Privileged EXEC Commands

**debug platform cli-redirection main**  
**debug platform configuration**  
**debug platform ipc**  
**debug platform led stack**

## FallBack Bridging

### Unsupported Privileged EXEC Commands

**clear bridge** [*bridge-group*] **multicast** [*router-ports* | **groups** | **counts**] [*group-address*] [*interface-unit*] [**counts**]  
**clear vlan statistics**  
**show bridge** [*bridge-group*] **circuit-group** [*circuit-group*] [*src-mac-address*] [*dst-mac-address*]  
**show bridge** [*bridge-group*] **multicast** [*router-ports* | **groups**] [*group-address*]  
**show bridge vlan**  
**show interfaces crb**  
**show interfaces** {**ethernet** | **fastethernet**} [*interface* | *slot/port*] **irb**  
**show subscriber-policy** *range*

### Unsupported Global Configuration Commands

**bridge** *bridge-group* **acquire**  
**bridge** *bridge-group* **address** *mac-address* {**forward** | **discard**} [*interface-id*]  
**bridge** *bridge-group* **aging-time** *seconds*  
**bridge** *bridge-group* **bitswap\_l3\_addresses**  
**bridge** *bridge-group* **bridge ip**  
**bridge** *bridge-group* **circuit-group** *circuit-group* **pause** *milliseconds*  
**bridge** *bridge-group* **circuit-group** *circuit-group* **source-based**

**bridge cmf**  
**bridge crb**  
**bridge** *bridge-group domain domain-name*  
**bridge irb**  
**bridge** *bridge-group mac-address-table limit number*  
**bridge** *bridge-group multicast-source*  
**bridge** *bridge-group protocol dec*  
**bridge** *bridge-group route protocol*  
**bridge** *bridge-group subscriber policy policy*  
**subscriber-policy** *policy* [[**no** | **default**] *packet* [**permit** | **deny**]]

## Unsupported Interface Configuration Commands

**bridge-group** *bridge-group cbus-bridging*  
**bridge-group** *bridge-group circuit-group circuit-number*  
**bridge-group** *bridge-group input-address-list access-list-number*  
**bridge-group** *bridge-group input-lat-service-deny group-list*  
**bridge-group** *bridge-group input-lat-service-permit group-list*  
**bridge-group** *bridge-group input-lsap-list access-list-number*  
**bridge-group** *bridge-group input-pattern-list access-list-number*  
**bridge-group** *bridge-group input-type-list access-list-number*  
**bridge-group** *bridge-group lat-compression*  
**bridge-group** *bridge-group output-address-list access-list-number*  
**bridge-group** *bridge-group output-lat-service-deny group-list*  
**bridge-group** *bridge-group output-lat-service-permit group-list*  
**bridge-group** *bridge-group output-lsap-list access-list-number*  
**bridge-group** *bridge-group output-pattern-list access-list-number*  
**bridge-group** *bridge-group output-type-list access-list-number*  
**bridge-group** *bridge-group sse*  
**bridge-group** *bridge-group subscriber-loop-control*  
**bridge-group** *bridge-group subscriber-trunk*  
**bridge** *bridge-group lat-service-filtering*  
**frame-relay map** **bridge** *dlci broadcast*  
**interface** **bvi** *bridge-group*  
**x25 map** **bridge** *x.121-address broadcast* [*options-keywords*]

# HSRP

## Unsupported Global Configuration Commands

`interface Async`  
`interface BVI`  
`interface Dialer`  
`interface Group-Async`  
`interface Lex`  
`interface Multilink`  
`interface Virtual-Template`  
`interface Virtual-Tokenring`

## Unsupported Interface Configuration Commands

`mtu`  
`standby mac-refresh seconds`  
`standby use-bia`

# IGMP Snooping

## Unsupported Global Configuration Commands

`ip igmp snooping report-suppression`  
`ip igmp snooping source-only-learning`  
`ip igmp snooping tcn`

# Interfaces

## Unsupported Privileged EXEC Commands

`show interfaces [interface-id | vlan vlan-id] [crb | fair-queue | irb | mac-accounting | precedence | irb | random-detect | rate-limit | shape]`

## Unsupported Global Configuration Commands

`interface tunnel`

## Unsupported Interface Configuration Commands

**bandwidth** {*bandwidth-kbps* | **percent** *percent*}

**random-detect** [**dscp-based** | **precedence-based**]

**random-detect dscp** *dscp min-threshold max-threshold mark-prob-denominator*

**random-detect exponential-weighting-constant** *weight*

**random-detect precedence** *ip-precedence min-threshold max-threshold mark-prob-denominator*

**switchport broadcast** *level*

**switchport multicast** *level*

**switchport unicast** *level*



### Note

The **switchport broadcast**, **switchport multicast**, and **switchport unicast** commands have been replaced by the **storm-control** {**broadcast** | **multicast** | **unicast**} **level** *level* [*.level*] interface configuration command.

## IP Multicast Routing

### Unsupported Privileged EXEC Commands

**clear ip rtp header-compression** [*type number*]

The **debug ip packet** command displays packets received by the switch CPU. It does not display packets that are hardware-switched.

The **debug ip mcache** command affects packets received by the switch CPU. It does not display packets that are hardware-switched.

The **debug ip mpacket** [**detail**] [*access-list-number* [*group-name-or-address*]] command affects only packets received by the switch CPU. Because most multicast packets are hardware-switched, use this command only when you know that the route will forward the packet to the CPU.

**debug ip pim atm**

**show frame-relay ip rtp header-compression** [**interface** *type number*]

The **show ip mcache** command displays entries in the cache for those packets that are sent to the switch CPU. Because most multicast packets are switched in hardware without CPU involvement, you can use this command, but multicast packet information is not displayed.

The **show ip mpacket** commands are supported but are only useful for packets received at the switch CPU. If the route is hardware-switched, the command has no effect because the CPU does not receive the packet and cannot display it.

**show ip pim vc** [*group-address* | *name*] [*type number*]

**show ip rtp header-compression** [*type number*] [**detail**]

## Unsupported Global Configuration Commands

**ip pim accept-rp** {*address* | **auto-rp**} [*group-access-list-number*]  
**ip pim message-interval** *seconds*

## Unsupported Interface Configuration Commands

**frame-relay ip rtp header-compression** [active | passive]  
**frame-relay map ip** *ip-address dlc* [**broadcast**] **compress**  
**frame-relay map ip** *ip-address dlc* **rtp header-compression** [active | passive]  
**ip igmp helper-address** *ip-address*  
**ip multicast helper-map** {*group-address* | **broadcast**} {*broadcast-address* | *multicast-address*}  
*extended-access-list-number*  
**ip multicast rate-limit** {in | out} [**video** | **whiteboard**] [**group-list** *access-list*] [**source-list** *access-list*]  
*kbps*  
**ip multicast ttl-threshold** *ttl-value* (instead, use the **ip multicast boundary** *access-list-number* interface configuration command)  
**ip multicast use-functional**  
**ip pim minimum-vc-rate** *pps*  
**ip pim multipoint-signalling**  
**ip pim nbma-mode**  
**ip pim vc-count** *number*  
**ip rtp compression-connections** *number*  
**ip rtp header-compression** [passive]

## IP Unicast Routing

### Unsupported Privileged EXEC or User EXEC Commands

**clear ip accounting** [checkpoint]  
**clear ip bgp** *address* **flap-statistics**  
**clear ip bgp** **prefix-list**  
**show cef** [drop | not-cef-switched]  
**show ip accounting** [checkpoint] [output-packets | access-violations]  
**show ip bgp** **dampened-paths**  
**show ip bgp** **inconsistent-as**  
**show ip bgp regexp** *regular expression*  
**show ip prefix-list** *regular expression*



## Unsupported Global Configuration Commands

**ip accounting-list** *ip-address wildcard*  
**ip as-path** access-list  
**ip accounting-transits** *count*  
**ip cef accounting** [**per-prefix**] [**non-recursive**]  
**ip cef traffic-statistics** [**load-interval** *seconds*] [**update-rate** *seconds*]  
**ip flow-aggregation**  
**ip flow-cache**  
**ip flow-export**  
**ip gratuitous-arps**  
**ip local**  
**ip prefix-list**  
**ip reflexive-list**  
**router** **egp**  
**router** **isis**  
**router** **iso-igrp**  
**router** **mobile**  
**router** **odr**  
**router** **static**

## Unsupported Interface Configuration Commands

**ip** **accounting**  
**ip** **load-sharing** [**per-packet**]  
**ip** **mtu** *bytes*  
**ip** **route-cache**  
**ip** **verify**  
**ip** **unnumbered** *type number*  
All **ip** **security** commands

## Unsupported BGP Router Configuration Commands

**address-family** vpnv4  
**default-information** originate  
**neighbor** advertise-map  
**neighbor** allowas-in  
**neighbor** default-originate  
**neighbor** description  
**network** backdoor  
**table-map**

## Unsupported Route Map Commands

**match** route-type  
**set as-path** {tag | **prepend** *as-path-string*}  
**set automatic-tag**  
**set dampening** *half-life reuse suppress max-suppress-time*  
**set default interface** *interface-id* [*interface-id*.....]  
**set interface** *interface-id* [*interface-id*.....]  
**set ip default next-hop** *ip-address* [*ip-address*.....]  
**set ip destination** *ip-address mask*  
**set ip precedence** *value*  
**set ip qos-group**  
**set metric-type** internal  
**set origin**  
**set metric-type** internal  
**set tag** *tag-value*

## Layer 2 Protocol Tunneling

### Unsupported Interface Configuration Commands

**l2protocol-tunnel** [**point-to-point** [**pagp** | **lACP** | **udld**]]

# Miscellaneous

## Unsupported Global Configuration Commands

`errdisable detect cause dhcp-rate-limit`  
`errdisable recovery cause dhcp-rate-limit`  
`errdisable recovery cause unicast flood`

## MPLS

All MPLS `traffic-eng` commands

## MSDP

## Unsupported Privileged EXEC Commands

`show access-expression`  
`show exception`  
`show location`  
`show pm LINE`  
`show smf [interface-id]`  
`show subscriber-policy [policy-number]`  
`show template [template-name]`

## Unsupported Global Configuration Commands

`ip msdp default-peer ip-address | name [prefix-list list]` (Because BGP/MBGP is not supported, use the `ip msdp peer` command instead of this command.)

## RADIUS

## Unsupported Global Configuration Commands

`aaa nas port extended`  
`radius-server attribute nas-port`  
`radius-server configure`  
`radius-server extended-portnames`

# SNMP

## Unsupported Global Configuration Commands

`snmp-server enable informs`  
`snmp-server enable traps flash insertion`  
`snmp-server enable traps flash removal`

# Spanning Tree

## Unsupported Global Configuration Commands

`spanning-tree etherchannel guard misconfig`  
`spanning-tree pathcost method {long | short}`

## Unsupported Interface Configuration Commands

`spanning-tree stack-port`

# Virtual Forwarding Infrastructure (VFI)

## Unsupported Global Configuration Commands

`l2 vfi vfi-name manual`  
All VFI configuration mode commands

## Unsupported Privileged EXEC Commands

`debug vfi`  
`show vfi`

# VLAN

## Unsupported vlan-config Commands

`private-vlan`

## Unsupported Privileged EXEC Commands

`show running-config vlan`

## Unsupported User EXEC Commands

`show vlan ifindex`

## VTP

## Unsupported Privileged EXEC Commands

`vtp {password password | pruning | version number} private-vlan`

**Note**

---

This command has been replaced by the **vtp** global configuration command.

---





---

## Numerics

802.1D

See STP

802.1p [12-1](#)

802.1Q

and trunk ports [9-3](#)

configuration limitations [10-19](#)

encapsulation [10-16](#)

native VLAN for untagged traffic [10-24](#)

tunneling

compatibility with other features [13-6](#)

defaults [13-4](#)

described [13-2](#)

tunnel ports with other features [13-6](#)

802.1s

See MSTP

802.1w

See RSTP

802.1x

See port-based authentication

802.3ad

See EtherChannel

802.3z flow control [9-15](#)

---

## A

abbreviating commands [2-4](#)

ABRs [28-28](#)

access-class command [25-17](#)

access control entries

See ACEs

access-denied response, VMPS [10-28](#)

access groups

applying ACLs to interfaces [25-18](#)

IP [25-19](#)

Layer 2 [25-18](#)

Layer 3 [25-19](#)

access lists

See ACLs

access ports

and Layer 2 protocol tunneling [13-12](#)

defined [9-2](#)

accounting

with RADIUS [7-28](#)

with TACACS+ [7-11, 7-16](#)

ACEs

and QoS [26-8](#)

defined [25-2](#)

Ethernet [25-2](#)

IP [25-2](#)

ACLs

ACEs [25-2](#)

any keyword [25-11](#)

applying

on bridged packets [25-36](#)

on multicast packets [25-37](#)

on routed packets [25-36](#)

on switched packets [25-35](#)

time ranges to [25-15](#)

to an interface [25-18](#)

to QoS [26-8](#)

classifying traffic for QoS [26-49](#)

comments in [25-17](#)

compiling [25-20](#)

configuring with VLAN maps [25-34](#)

## ACLs (continued)

- defined [25-1, 25-6](#)
- examples of [25-20, 26-49](#)
- extended IP
  - configuring for QoS classification [26-50](#)
  - creating [25-9](#)
  - matching criteria [25-6](#)
- hardware and software handling [25-19](#)
- host keyword [25-11](#)
- IP
  - applying to interfaces [25-18](#)
  - creating [25-6](#)
  - fragments and QoS guidelines [26-40](#)
  - implicit deny [25-8, 25-12, 25-14](#)
  - implicit masks [25-8](#)
  - matching criteria [25-6](#)
  - named [25-13](#)
  - terminal lines, setting on [25-17](#)
  - undefined [25-19](#)
  - violations, logging [25-14](#)
- limiting actions [25-35](#)
- logging messages [25-9](#)
- log keyword [25-14](#)
- MAC extended [25-25, 26-51](#)
- matching [25-6, 25-19](#)
- monitoring [25-38](#)
- named [25-13](#)
- number per QoS class map [26-40](#)
- numbers [25-7](#)
- port [25-2](#)
- precedence of [25-2](#)
- QoS [26-8, 26-49](#)
- router [25-2](#)
- standard IP
  - configuring for QoS classification [26-49](#)
  - creating [25-8](#)
  - matching criteria [25-6](#)
- supported features [25-19](#)
- time ranges [25-15](#)

## ACLs (continued)

- unsupported features [25-6](#)
- using router ACLs with VLAN maps [25-34](#)
- VLAN maps
  - configuration guidelines [25-28](#)
  - configuring [25-27](#)
- active router [29-1](#)
- address aliasing [17-2](#)
- addresses
  - displaying the MAC address table [5-25](#)
  - dynamic
    - accelerated aging [14-8](#)
    - changing the aging time [5-22](#)
    - default aging [14-8](#)
    - defined [5-20](#)
    - learning [5-21](#)
    - removing [5-22](#)
  - MAC, discovering [5-26](#)
  - multicast
    - group address range [31-3](#)
    - STP address management [14-8](#)
  - static
    - adding and removing [5-24](#)
    - defined [5-20](#)
- address resolution [5-26, 28-7](#)
- Address Resolution Protocol
  - See ARP
- adjacency tables, with CEF [28-87](#)
- administrative distances
  - defined [28-98](#)
  - OSPF [28-33](#)
  - routing protocol defaults [28-89](#)
- advertisements
  - CDP [19-1](#)
  - IGRP [28-23](#)
  - RIP [28-18](#)
  - VTP [10-20, 11-3](#)
- aggregate addresses, BGP [28-58](#)



- aggregated ports
    - See EtherChannel
  - aggregate policers [26-57](#)
  - aging, accelerating [14-8](#)
  - aging time
    - accelerated
      - for MSTP [15-20](#)
      - for STP [14-8, 14-21](#)
    - MAC address table [5-22](#)
    - maximum
      - for MSTP [15-21](#)
      - for STP [14-21](#)
  - alarms, RMON [22-3](#)
  - allowed-VLAN list [10-22](#)
  - alternate routes, IGRP [28-24](#)
  - Any Transport over MPLS
    - See AToM
  - area border routers
    - See ABRs
  - area routing
    - IS-IS [28-63](#)
    - ISO IGRP [28-63](#)
  - ARP
    - configuring [28-8](#)
    - defined [5-26, 28-8](#)
    - encapsulation [28-9](#)
    - static cache configuration [28-8](#)
  - ARP table
    - address resolution [5-26](#)
    - managing [5-26](#)
  - ASBRs [28-28](#)
  - AS-path filters, BGP [28-53](#)
  - asymmetrical links, and 802.1Q tunneling [13-4](#)
  - AToM [30-12](#)
  - attributes, RADIUS
    - vendor-proprietary [7-30](#)
    - vendor-specific [7-29](#)
  - audience [xxxiii](#)
  - authentication
    - EIGRP [28-41](#)
    - HSRP [29-8](#)
    - Kerberos [7-31](#)
    - local mode with AAA [7-36](#)
    - NTP associations [5-5](#)
    - RADIUS
      - key [7-21](#)
      - login [7-23](#)
    - See also port-based authentication
    - TACACS+
      - defined [7-11](#)
      - key [7-13](#)
      - login [7-14](#)
  - authentication keys, and routing protocols [28-99](#)
  - authoritative time source, described [5-2](#)
  - authorization
    - with RADIUS [7-27](#)
    - with TACACS+ [7-11, 7-16](#)
  - authorized ports with 802.1x [8-4](#)
  - autoconfiguration [3-3](#)
  - automatic QoS
    - See QoS
  - autonegotiation
    - interface configuration guidelines [9-13](#)
    - mismatches [34-7](#)
  - autonomous system boundary routers
    - See ASBRs
  - autonomous systems, in BGP [28-47](#)
  - Auto-RP, described [31-5](#)
  - auxiliary VLAN
    - See voice VLAN
  - availability features [1-3](#)
- 
- ## B
- BackboneFast
    - described [16-5](#)
    - enabling [16-14](#)

## banners

- configuring
  - login [5-20](#)
  - message-of-the-day login [5-19](#)
- default configuration [5-18](#)
- when displayed [5-18](#)

## BGP

- aggregate addresses [28-58](#)
- aggregate routes, configuring [28-58](#)
- CIDR [28-58](#)
- clear commands [28-62](#)
- community filtering [28-55](#)
- configuring neighbors [28-57](#)
- configuring routing sessions [30-9](#)
- default configuration [28-45](#)
- described [28-44](#)
- enabling [28-47](#)
- monitoring [28-62](#)
- multipath support [28-50](#)
- neighbors, types of [28-47](#)
- path selection [28-50](#)
- peers, configuring [28-57](#)
- prefix filtering [28-54](#)
- resetting sessions [28-49](#)
- route dampening [28-61](#)
- route maps [28-52](#)
- route reflectors [28-60](#)
- routing domain confederation [28-59](#)
- routing session with multi-VRF CE [28-81](#)
- show commands [28-62](#)
- supernets [28-58](#)
- Version 4 [28-44](#)

blocking packets [18-5](#)

## booting

- boot loader, function of [3-2](#)
- boot process [3-1](#)
- manually [3-12](#)
- specific image [3-12](#)

## boot loader

- accessing [3-13](#)
- described [3-2](#)
- environment variables [3-13](#)
- prompt [3-13](#)
- trap-door mechanism [3-2](#)

bootstrap router (BSR), described [31-5](#)

## Border Gateway Protocol

See BGP

## BPDU

- error-disabled state [16-3](#)
- filtering [16-3](#)
- RSTP format [15-9](#)

## BPDU filtering

- described [16-3](#)
- enabling [16-12](#)

## BPDU guard

- described [16-3](#)
- enabling [16-11](#)

bridged packets, ACLs on [25-36](#)

## bridge groups

See fallback bridging

## bridge protocol data unit

See BPDU

broadcast flooding [28-15](#)

## broadcast packets

- directed [28-12](#)
- flooded [28-12](#)

broadcast storm-control command [18-3](#)broadcast storms [18-1, 28-12](#)bypass mode, Layer 2 protocol-tunneling [13-11, 13-13](#)


---

**C**
cables, monitoring for unidirectional links [20-1](#)caution, described [xxxiv](#)

**CBWFQ**

- configuring
  - with DSCP-based WRED [26-89](#)
  - with IP precedence-based WRED [26-93](#)
  - with tail drop [26-86](#)
- described [26-28](#)

**CDP**

- and trusted boundary [26-46](#)
- configuring [19-2](#)
- default configuration [19-2](#)
- described [19-1](#)
- disabling for routing device [19-3, 19-4](#)
- enabling and disabling
  - on an interface [19-4](#)
  - on a switch [19-3](#)
- Layer 2 protocol tunneling [13-10](#)
- monitoring [19-5](#)
- overview [19-1](#)
- transmission timer and holdtime, setting [19-2](#)
- updates [19-2](#)

**CEF** [28-86](#)**CGMP**

- as IGMP snooping learning method [17-7](#)
- clearing cached group entries [31-49](#)
- enabling server support [31-32](#)
- joining multicast group [17-2](#)
- overview [31-7](#)
- server support only [31-7](#)

**CIDR** [28-58](#)**Cisco 7960 IP Phone** [12-1](#)**Cisco Discovery Protocol**

- See CDP

**Cisco Express Forwarding**

- See CEF

**Cisco Group Management Protocol**

- See CGMP

**Cisco Intelligence Engine 2100 Series Configuration Registrar**

- See IE2100

**Cisco IOS File System** [1-3](#)**Cisco Networking Services**

- See IE2100

**CiscoWorks 2000** [1-2, 24-4](#)**class-based weighted fair queueing**

- See CBWFQ

**classless interdomain routing**

- See CIDR

**classless routing** [28-6](#)**class maps for QoS**

- configuring
  - egress [26-78](#)
  - ingress [26-52](#)
- described [26-8, 26-23](#)
- displaying [26-75, 26-101](#)

**class of service**

- See CoS

**clearing interfaces** [9-21](#)**CLI**

- abbreviating commands [2-4](#)
- command modes [2-1](#)
- editing features
  - enabling and disabling [2-7](#)
  - keystroke editing [2-7](#)
  - wrapped lines [2-8](#)
- error messages [2-5](#)
- filtering command output [2-9](#)
- getting help [2-3](#)
- history
  - changing the buffer size [2-5](#)
  - described [2-5](#)
  - disabling [2-6](#)
  - recalling commands [2-6](#)
  - no and default forms of commands [2-4](#)
- client mode, VTP [11-3](#)
- CLNS
  - See ISO CLNS
- clock
  - See system clock

- command-line interface
  - See CLI
- command modes [2-1](#)
- commands
  - abbreviating [2-4](#)
  - no and default [2-4](#)
  - setting privilege levels [7-8](#)
- community list, BGP [28-56](#)
- community strings
  - configuring [24-7](#)
  - overview [24-4](#)
- config.text [3-11](#)
- configuration, initial
  - Express Setup [1-8](#)
  - See also hardware installation guide
  - setup (CLI) program [1-8](#)
- configuration examples
  - ACLs and VLAN maps [25-29](#)
  - IP ACLs [25-20](#)
  - multi-VRF-CE [28-82](#)
  - network [1-11](#)
  - SNMP [24-14](#)
  - VLAN maps [25-31](#)
  - VMPS [10-29](#)
- configuration files
  - clearing the startup configuration [B-18](#)
  - creating using a text editor [B-9](#)
  - default name [3-11](#)
  - deleting a stored configuration [B-19](#)
  - described [B-8](#)
  - downloading
    - automatically [3-11](#)
    - preparing [B-10, B-12, B-16](#)
    - reasons for [B-8](#)
    - using FTP [B-13](#)
    - using RCP [B-16](#)
    - using TFTP [B-11](#)
  - guidelines for creating and using [B-8](#)
  - invalid combinations when copying [B-5](#)
  - configuration files (continued)
    - limiting TFTP server access [24-13](#)
    - obtaining with DHCP [3-7](#)
    - password recovery disable considerations [7-5](#)
    - specifying the filename [3-11](#)
    - system contact and location information [24-13](#)
    - types and location [B-9](#)
    - uploading
      - preparing [B-10, B-12, B-16](#)
      - reasons for [B-8](#)
      - using FTP [B-14](#)
      - using RCP [B-17](#)
      - using TFTP [B-11](#)
- configuration guidelines
  - 802.1Q trunks [10-19](#)
  - applying ACLs [25-18](#)
  - applying MAC ACLs [25-26](#)
  - auto-QoS [26-33](#)
  - Auto-RP and BSR [31-9](#)
  - EtherChannel [27-9](#)
  - fallback bridging [33-3](#)
  - hierarchical QoS [26-76](#)
  - HSRP
    - authentication and timers [29-8](#)
    - interfaces [29-4](#)
    - priority [29-6](#)
  - interface speed and duplex mode [9-13](#)
  - IP multicast routing [31-8](#)
  - MSTP [15-12, 16-9](#)
  - multi-VRF CE [28-79](#)
  - named ACLs [25-13](#)
  - port-based authentication [8-11](#)
  - port security [18-9](#)
  - protected ports [18-5](#)
  - router ACLs and VLAN maps [25-34](#)
  - RSPAN [21-16](#)
  - SDM templates [6-3](#)
  - SNMP [24-6](#)
  - SPAN [21-10](#)

- configuration guidelines (continued)
  - standard QoS [26-40](#)
  - STP [14-12, 16-9](#)
  - UDLD [20-4](#)
  - VLAN maps [25-28](#)
  - VLANs
    - extended-range [10-13](#)
    - normal-range [10-6](#)
  - VMPS [10-29](#)
  - voice VLAN [12-3](#)
  - VTP [11-8](#)
- configuration settings, saving [3-10](#)
- configure terminal command [9-7](#)
- config-vlan mode [2-2, 10-7](#)
- congestion-avoidance mechanisms
  - WRED [26-27, 26-89, 26-93](#)
  - WTD [26-13, 26-65, 26-69](#)
- congestion management for QoS [26-27](#)
- connectionless service, and VPNs [30-3](#)
- connections, secure remote [7-37](#)
- connectivity problems [34-8, 34-10, 34-11](#)
- consistency checks in VTP version 2 [11-4](#)
- console port, connecting to [2-10](#)
- conventions
  - command [xxxiv](#)
  - for examples [xxxiv](#)
  - publication [xxxiv](#)
  - text [xxxiv](#)
- corrupted software, recovery steps with XMODEM [34-2](#)
- CoS
  - configuring the default port value [26-45](#)
  - in Layer 2 frames [26-2](#)
  - override priority [12-5](#)
  - trust priority [12-5](#)
- CoS input queue threshold map for QoS [26-16](#)
- CoS output queue threshold map for QoS [26-19](#)
- CoS-to-DSCP map for QoS [26-59](#)
- counters, clearing interface [9-21](#)
- crashinfo file [34-17](#)

- cryptographic software image
  - Kerberos [7-31](#)
  - SSH [7-37](#)
- customer edge devices
  - and Multi-VFR CE [28-76](#)
  - and VPNs [30-3, 30-5](#)

---

## D

- daylight saving time [5-13](#)
- debugging
  - enabling all system diagnostics [34-14](#)
  - enabling for a specific feature [34-13](#)
  - redirecting error message output [34-14](#)
  - using commands [34-13](#)
- default commands [2-4](#)
- default configuration
  - 802.1Q tunneling [13-4](#)
  - 802.1x [8-10](#)
  - auto-QoS [26-30](#)
  - banners [5-18](#)
  - BGP [28-45](#)
  - booting [3-11](#)
  - CDP [19-2](#)
  - DNS [5-17](#)
  - EIGRP [28-38](#)
  - EoMPLS [30-15](#)
  - EtherChannel [27-9](#)
  - fallback bridging [33-3](#)
  - hierarchical QoS [26-76](#)
  - HSRP [29-4](#)
  - IGMP [31-26](#)
  - IGMP filtering [17-19](#)
  - IGMP snooping [17-5](#)
  - IGRP [28-24](#)
  - initial switch information [3-3](#)
  - IP addressing, IP routing [28-4](#)
  - IP multicast routing [31-8](#)
  - IS-IS [28-67](#)

## default configuration (continued)

- ISO IGRP [28-64](#)
- Layer 2 interfaces [9-11](#)
- Layer 2 protocol tunneling [13-13](#)
- MAC address table [5-21](#)
- MPLS [30-6](#)
- MPLS QoS [30-20](#)
- MSDP [32-4](#)
- MSTP [15-12](#)
- multi-VRF CE [28-78](#)
- MVR [17-14](#)
- NTP [5-4](#)
- optional spanning-tree features [16-9](#)
- OSPF [28-29](#)
- password and privilege level [7-2](#)
- RADIUS [7-20](#)
- RIP [28-18](#)
- RMON [22-3](#)
- RSPAN [21-9](#)
- SNMP [24-6](#)
- SPAN [21-9](#)
- standard QoS [26-38](#)
- STP [14-11](#)
- system message logging [23-3](#)
- system name and prompt [5-15](#)
- TACACS+ [7-13](#)
- UDLD [20-4](#)
- VLAN, Layer 2 Ethernet interfaces [10-19](#)
- VLAN mapping [13-8](#)
- VLANs [10-8](#)
- VMPS [10-29](#)
- voice VLAN [12-3](#)
- VTP [11-6](#)

default gateway [3-10, 28-10](#)

default networks [28-89](#)

default routes [28-89](#)

default routing [28-2](#)

deleting VLANs [10-11](#)

description command [9-17](#)

designing your network, examples [1-11](#)

destination addresses, in ACLs [25-10](#)

destination-IP address based forwarding,  
EtherChannel [27-7](#)

destination-MAC address forwarding, EtherChannel [27-7](#)

detecting indirect link failures, STP [16-6](#)

device discovery protocol [19-1](#)

DHCP-based autoconfiguration

- client request message exchange [3-4](#)
- configuring
  - client side [3-3](#)
  - DNS [3-6](#)
  - relay device [3-6](#)
  - server-side [3-5](#)
  - TFTP server [3-5](#)
- example [3-8](#)
- lease options
  - for IP address information [3-5](#)
  - for receiving the configuration file [3-5](#)
- overview [3-3](#)
- relationship to BOOTP [3-3](#)

Differentiated Services architecture, QoS [26-2](#)

Differentiated Services Code Point

- See DSCP

Diffusing Update Algorithm (DUAL) [28-37](#)

directories

- changing [B-3](#)
- creating and removing [B-4](#)
- displaying the working [B-3](#)

Distance Vector Multicast Routing Protocol

- See DVMRP

distance-vector protocols [28-2](#)

distribute-list command [28-97](#)

DNS

- and DHCP-based autoconfiguration [3-6](#)
- default configuration [5-17](#)
- displaying the configuration [5-18](#)
- overview [5-16](#)
- setting up [5-17](#)

- documentation
  - feedback [xxxvi](#)
  - obtaining, world wide web [xxxv](#)
  - ordering [xxxvi](#)
  - related [xxxv](#)
- document conventions [xxxiv](#)
- domain names
  - DNS [5-16](#)
  - VTP [11-8](#)
- Domain Name System
  - See DNS
- domains, ISO IGRP routing [28-63](#)
- dot1q-tunnel switchport mode [10-18](#)
- double-tagged packets
  - 802.1Q tunneling [13-2](#)
  - Layer 2 protocol tunneling [13-12](#)
- downloading
  - configuration files
    - preparing [B-10, B-12, B-16](#)
    - reasons for [B-8](#)
    - using FTP [B-13](#)
    - using RCP [B-16](#)
    - using TFTP [B-11](#)
  - image files
    - deleting old image [B-23](#)
    - preparing [B-21, B-24, B-28](#)
    - reasons for [B-19](#)
    - using FTP [B-25](#)
    - using RCP [B-29](#)
    - using TFTP [B-22](#)
- drop threshold for Layer 2 protocol packets [13-13](#)
- DSCP [26-2](#)
- DSCP input queue threshold map for QoS [26-16](#)
- DSCP output queue threshold map for QoS [26-19](#)
- DSCP-to-CoS map for QoS [26-62](#)
- DSCP-to-DSCP-mutation map for QoS [26-63](#)
- DTP [10-17](#)
- DUAL finite state machine, EIGRP [28-37](#)
- duplex mode, configuring [9-12](#)
- DVMRP
  - autosummarization
    - configuring a summary address [31-46](#)
    - disabling [31-48](#)
  - connecting PIM domain to DVMRP router [31-38](#)
  - enabling unicast routing [31-42](#)
  - interoperability
    - with Cisco devices [31-36](#)
    - with IOS software [31-7](#)
  - mrinfo requests, responding to [31-41](#)
  - neighbors
    - advertising the default route to [31-40](#)
    - discovery with Probe messages [31-36](#)
    - displaying information [31-41](#)
    - prevent peering with nonpruning [31-44](#)
    - rejecting nonpruning [31-43](#)
  - overview [31-7](#)
  - routes
    - adding a metric offset [31-48](#)
    - advertising all [31-48](#)
    - advertising the default route to neighbors [31-40](#)
    - caching DVMRP routes learned in report messages [31-42](#)
    - changing the threshold for syslog messages [31-45](#)
    - deleting [31-49](#)
    - displaying [31-50](#)
    - favoring one over another [31-48](#)
    - limiting the number injected into MBONE [31-45](#)
    - limiting unicast route advertisements [31-36](#)
  - routing table [31-7](#)
  - source distribution tree, building [31-7](#)
  - tunnels
    - configuring [31-38](#)
    - displaying neighbor information [31-41](#)

dynamic access ports  
     characteristics [10-4](#)  
     configuring [10-31](#)  
     defined [9-3](#)

dynamic addresses  
     See addresses

dynamic auto trunking mode [10-18](#)

dynamic desirable trunking mode [10-18](#)

Dynamic Host Configuration Protocol  
     See DHCP-based autoconfiguration

dynamic port VLAN membership  
     described [10-29](#)  
     reconfirming [10-31, 10-32](#)  
     troubleshooting [10-33](#)  
     types of connections [10-31](#)

dynamic routing  
     ISO CLNS [28-63](#)  
     protocols [28-2](#)

Dynamic Trunking Protocol  
     See DTP

## E

EBGP [28-43, 30-5](#)

editing features  
     enabling and disabling [2-7](#)  
     keystrokes used [2-7](#)  
     wrapped lines [2-8](#)

### EIGRP

    and IGRP [28-39](#)  
     authentication [28-41](#)  
     components [28-37](#)  
     configuring [28-39](#)  
     default configuration [28-38](#)  
     definition [28-37](#)  
     interface parameters, configuring [28-40](#)  
     monitoring [28-42](#)  
     enable password [7-3](#)

    enable secret password [7-3](#)

    encryption for passwords [7-3](#)

Enhanced IGRP

    See EIGRP

enhanced-services interfaces

    See ES interfaces

environment variables, function of [3-14](#)

### EoMPLS

    and 802.1Q tunneling [30-13](#)  
     and Layer 2 protocol tunneling [30-14](#)  
     and QoS [30-14](#)  
     configuring [30-16](#)  
     default configuration [30-15](#)  
     limitations [30-14](#)  
     monitoring [30-22](#)  
     packet flow [30-17](#)  
     QoS [30-18](#)

equal-cost routing [28-87](#)

error messages during command entry [2-5](#)

ES interfaces [9-7, 30-1](#)

### EtherChannel

    802.3ad, described [27-5](#)  
     automatic creation of [27-4, 27-5](#)  
     channel groups  
         binding physical and logical interfaces [27-3](#)  
         numbering of [27-3](#)  
     configuration guidelines [27-9](#)  
     configuring  
         Layer 3 physical interfaces [27-13](#)  
         Layer 3 port-channel logical interfaces [27-12](#)  
     configuring Layer 2 interfaces [27-10](#)  
     default configuration [27-9](#)  
     described [27-2](#)  
     displaying status [27-19](#)  
     forwarding methods [27-6, 27-15](#)  
     interaction  
         with STP [27-9](#)  
         with VLANs [27-10](#)



## EtherChannel (continued)

## LACP

- described [27-5](#)
- displaying status [27-19](#)
- hot-standby ports [27-17](#)
- interaction with other features [27-6](#)
- modes [27-6](#)
- port priority [27-18](#)
- system priority [27-18](#)

Layer 3 interface [28-3](#)load balancing [27-6, 27-15](#)logical interfaces, described [27-3](#)number of interfaces per [27-2](#)

## PAgP

- aggregate-port learners [27-16](#)
- compatibility with Catalyst 1900 [27-16](#)
- described [27-4](#)
- displaying status [27-19](#)
- interaction with other features [27-5](#)
- learn method and priority configuration [27-16](#)
- modes [27-4](#)
- silent mode [27-5](#)

## port-channel interfaces

- described [27-3](#)
- numbering of [27-3](#)

port groups [9-5](#)

## Ethernet over MPLS

## See EoMPLS

## Ethernet VLANs

- creating [10-9](#)
- defaults and ranges [10-8](#)
- modifying [10-9](#)

events, RMON [22-3](#)

## examples

- conventions for [xxxiv](#)
- network configuration [1-11](#)

experimental bits, setting MPLS priority with [30-20](#)

## Express Setup

See hardware installation guide

## extended-range VLANs

- configuration guidelines [10-13](#)
- configuring [10-12](#)
- creating [10-13, 10-14](#)
- defined [10-1](#)

## extended system ID

- MSTP [15-14](#)
- STP [14-4, 14-14](#)

Extensible Authentication Protocol over LAN [8-1](#)exterior routes, IGRP [28-23](#)

## external BGP

See EBGp

external neighbors, BGP [28-47](#)


---

**F**

## fallback bridging

and protected ports [33-3](#)

## bridge groups

- creating [33-3](#)
- described [33-1](#)
- displaying [33-10](#)
- function of [33-2](#)
- number supported [33-4](#)
- removing [33-4](#)

## bridge table

- clearing [33-10](#)
- displaying [33-10](#)

configuration guidelines [33-3](#)connecting interfaces with [9-6](#)default configuration [33-3](#)described [33-1](#)

## frame forwarding

- flooding packets [33-2](#)
- forwarding packets [33-2](#)

overview [33-1](#)

## fallback bridging (continued)

## STP

- disabling on an interface [33-10](#)
- forward-delay interval [33-9](#)
- hello BPDU interval [33-8](#)
- interface priority [33-6](#)
- maximum-idle interval [33-9](#)
- path cost [33-7](#)
- VLAN-bridge spanning-tree priority [33-5](#)
- VLAN-bridge STP [33-2](#)
- SVIs and routed ports [33-1](#)
- VLAN-bridge STP [14-11](#), [33-1](#)

## features

- availability [1-3](#)
- Layer 2 VPN services [1-4](#)
- Layer 3 [1-7](#)
- Layer 3 VPN services [1-5](#)
- manageability [1-3](#)
- management options [1-2](#)
- monitoring [1-8](#)
- performance [1-2](#)
- QoS [1-6](#)
- security [1-5](#)
- VLAN [1-4](#)

feedback to Cisco Systems, web [xxxvi](#)

FIB [28-86](#)

fiber-optic, detecting unidirectional links [20-1](#)

## files

- copying [B-4](#)
- crashinfo
  - description [34-17](#)
  - displaying the contents of [34-17](#)
  - location [34-17](#)
- deleting [B-5](#)
- displaying the contents of [B-7](#)

## files (continued)

## tar

- creating [B-6](#)
- displaying the contents of [B-6](#)
- extracting [B-7](#)
- image file format [B-20](#)

## file system

- displaying available file systems [B-2](#)
- displaying file information [B-3](#)
- local file system names [B-1](#)
- network file system names [B-4](#)
- setting the default [B-3](#)

## filtering

- in a VLAN [25-27](#)
- non-IP traffic [25-25](#)
- show and more command output [2-9](#)

## filters, IP

See ACLs, IP

flash device, number of [B-1](#)

flash updates, IGRP [28-24](#)

flooded traffic, blocking [18-6](#)

## flowcharts

- QoS egress, two-rate policing and marking [26-24](#)
- QoS egress hierarchical queues queueing and scheduling [26-26](#)
- QoS egress queue-set queueing and scheduling [26-17](#)
- QoS ingress, single-rate policing and marking [26-11](#)
- QoS ingress classification [26-7](#)
- QoS ingress queueing and scheduling [26-15](#)

flow control [9-15](#)

## forward-delay time

- MSTP [15-20](#)
- STP [14-21](#)

forwarding equivalence classes [30-2](#)

## Forwarding Information Base

See FIB

forwarding non-routable protocols [33-1](#)

## FTP

accessing MIB files [A-3](#)

configuration files

downloading [B-13](#)

overview [B-12](#)

preparing the server [B-12](#)

uploading [B-14](#)

image files

deleting old image [B-27](#)

downloading [B-25](#)

preparing the server [B-24](#)

uploading [B-27](#)

## G

get-bulk-request operation [24-3](#)

get-next-request operation [24-3, 24-4](#)

get-request operation [24-3, 24-4](#)

get-response operation [24-3](#)

global configuration mode [2-2](#)

guest VLAN and 802.1x [8-8](#)

guide

audience [xxxiii](#)

purpose of [xxxiii](#)

## H

hardware limitations and Layer 3 interfaces [9-18](#)

hello time

MSTP [15-19](#)

STP [14-20](#)

help, for the command line [2-3](#)

hierarchical QoS

See QoS

history

changing the buffer size [2-5](#)

described [2-5](#)

disabling [2-6](#)

recalling commands [2-6](#)

history table, level and number of syslog messages [23-9](#)

hosts, limit on dynamic ports [10-33](#)

Hot Standby Router Protocol

See HSRP

HP OpenView [1-2](#)

HSRP

authentication string [29-8](#)

configuring [29-3](#)

default configuration [29-4](#)

definition [29-1](#)

enabling [29-5](#)

guidelines [29-4](#)

monitoring [29-10](#)

overview [29-1](#)

priority [29-6](#)

timers [29-8](#)

tracking [29-6](#)

IBGP [28-43, 30-5](#)

ICMP

redirect messages [28-10](#)

time exceeded messages [34-12](#)

traceroute and [34-12](#)

unreachable messages [25-18](#)

unreachables and ACLs [25-19](#)

ICMP ping

executing [34-9](#)

overview [34-8](#)

ICMP Router Discovery Protocol

See IRDP

## IDS appliances

and ingress RSPAN [21-20](#)and ingress SPAN [21-13](#)

## IE2100

## CNS embedded agents

described [4-5](#)enabling automated configuration [4-6](#)enabling configuration agent [4-9](#)enabling event agent [4-8](#)

## Configuration Registrar

configID, deviceID, hostname [4-3](#)configuration service [4-2](#)described [4-1](#)event service [4-3](#)

## IGMP

## configuring the switch

as a member of a group [31-26](#)statically connected member [31-31](#)controlling access to groups [31-27](#)default configuration [31-26](#)deleting cache entries [31-50](#)displaying groups [31-50](#)fast switching [31-31](#)host-query interval, modifying [31-29](#)joining multicast group [17-2](#)join messages [17-2](#)leave processing, enabling [17-10](#)leaving multicast group [17-4](#)multicast reachability [31-26](#)overview [31-2](#)queries [17-3](#)

## Version 1

changing to Version 2 [31-28](#)described [31-3](#)

## IGMP (continued)

## Version 2

changing to Version 1 [31-28](#)described [31-3](#)maximum query response time value [31-30](#)pruning groups [31-30](#)query timeout value [31-30](#)

## IGMP filtering

configuring [17-19](#)default configuration [17-19](#)described [17-19](#)monitoring [17-22](#)IGMP groups, setting the maximum number [17-21](#)

## IGMP profile

applying [17-20](#)configuration mode [17-19](#)configuring [17-20](#)

## IGMP snooping

and address aliasing [17-2](#)configuring [17-5](#)default configuration [17-5](#)definition [17-2](#)enabling and disabling [17-6](#)global configuration [17-6](#)Immediate Leave [17-5](#)method [17-6](#)monitoring [17-10](#)VLAN configuration [17-6](#)IGP [28-28](#)

## IGRP

advertisements [28-23](#)alternate routes [28-24](#)configuring [28-25](#)default configuration [28-24](#)described [28-23](#)exterior routes [28-23](#)flash updates [28-24](#)interior routes [28-23](#)load balancing [28-24](#)

- IGRP (continued)
  - poison-reverse updates [28-24](#)
  - split horizon [28-27](#)
  - system routes [28-23](#)
  - traffic sharing [28-25](#)
  - unequal-cost load balancing [28-24](#)
- Immediate-Leave, IGMP [17-5](#)
- initial configuration
  - defaults [1-8](#)
  - Express Setup [1-8](#)
  - See also hardware installation guide
  - setup (CLI) program [1-8](#)
- Intelligence Engine 2100 Series CNS Agents
  - See IE2100
- interdomain routing, ISO IGRP [28-63](#)
- interface
  - number [9-7](#)
  - range macros [9-9](#)
- interface command [9-6, 9-7](#)
- interface configuration mode [2-3](#)
- interfaces
  - configuration guidelines [9-13](#)
  - configuring [9-7](#)
  - configuring duplex mode [9-12](#)
  - configuring speed [9-12](#)
  - counters, clearing [9-21](#)
  - described [9-17](#)
  - descriptive name, adding [9-17](#)
  - displaying information about [9-20](#)
  - flow control [9-15](#)
  - management [1-2](#)
  - monitoring [9-20](#)
  - naming [9-17](#)
  - physical, identifying [9-6](#)
  - range of [9-8](#)
  - restarting [9-22](#)
  - shutting down [9-22](#)
  - supported [9-6](#)
  - types of [9-1](#)
- interfaces range macro command [9-9](#)
- interface types [9-6](#)
- Interior Gateway Protocol
  - See IGP
- Interior Gateway Routing Protocol
  - See IGRP
- interior routes, IGRP [28-23](#)
- Intermediate-System-to-Intermediate-System protocol
  - See IS-IS
- internal BGP
  - See IBGP
- internal neighbors, BGP [28-47](#)
- Internet Control Message Protocol
  - See ICMP
- Internet Group Management Protocol
  - See IGMP
- Inter-Switch Link
  - See ISL
- inter-VLAN routing [28-2](#)
- Intrusion Detection System
  - See IDS appliances
- ip access group command [25-19](#)
- IP ACLs
  - extended, creating [25-9](#)
  - for QoS classification [26-8](#)
  - implicit deny [25-8, 25-12, 25-14](#)
  - implicit masks [25-8](#)
  - logging [25-14](#)
  - named [25-13](#)
  - standard, creating [25-8](#)
  - undefined [25-19](#)
  - virtual terminal lines, setting on [25-17](#)
- IP addresses
  - classes of [28-5](#)
  - default configuration [28-4](#)
  - discovering [5-26](#)
  - for IP routing [28-4](#)
  - MAC address association [28-7](#)
  - monitoring [28-16](#)

- IP broadcast address [28-15](#)
- ip cef command [28-87](#)
- IP directed broadcasts [28-13](#)
- ip igmp profile command [17-19](#)
- IP information
  - assigned
    - manually [3-9](#)
    - through DHCP-based autoconfiguration [3-3](#)
  - default configuration [3-3](#)
- IP multicast routing
  - addresses
    - all-hosts [31-3](#)
    - all-multicast-routers [31-3](#)
    - host group address range [31-3](#)
  - administratively-scoped boundaries, described [31-34](#)
  - and IGMP snooping [17-2](#)
  - Auto-RP
    - adding to an existing sparse-mode cloud [31-14](#)
    - benefits of [31-13](#)
    - clearing the cache [31-50](#)
    - configuration guidelines [31-9](#)
    - filtering incoming RP announcement messages [31-16](#)
    - overview [31-5](#)
    - preventing candidate RP spoofing [31-16](#)
    - preventing join messages to false RPs [31-15](#)
    - setting up in a new internetwork [31-13](#)
    - using with BSR [31-21](#)
  - bootstrap router
    - configuration guidelines [31-9](#)
    - configuring candidate BSRs [31-19](#)
    - configuring candidate RPs [31-20](#)
    - defining the IP multicast boundary [31-18](#)
    - defining the PIM domain border [31-17](#)
    - overview [31-5](#)
    - using with Auto-RP [31-21](#)
  - Cisco implementation [31-2](#)
  - configuring
    - basic multicast routing [31-10](#)
    - IP multicast boundary [31-34](#)
- IP multicast routing (continued)
  - default configuration [31-8](#)
  - enabling
    - multicast forwarding [31-10](#)
    - PIM mode [31-11](#)
  - group-to-RP mappings
    - Auto-RP [31-5](#)
    - BSR [31-5](#)
  - MBONE
    - deleting sdr cache entries [31-50](#)
    - described [31-33](#)
    - displaying sdr cache [31-51](#)
    - enabling sdr listener support [31-33](#)
    - limiting DVMRP routes advertised [31-45](#)
    - limiting sdr cache entry lifetime [31-34](#)
    - SAP packets for conference session announcement [31-33](#)
    - Session Directory (sdr) tool, described [31-33](#)
  - monitoring
    - packet rate loss [31-51](#)
    - peering devices [31-51](#)
    - tracing a path [31-51](#)
  - multicast forwarding, described [31-6](#)
  - PIMv1 and PIMv2 interoperability [31-9](#)
  - protocol interaction [31-2](#)
  - reverse path check (RPF) [31-6](#)
  - routing table
    - deleting [31-50](#)
    - displaying [31-50](#)
  - RP
    - assigning manually [31-12](#)
    - configuring Auto-RP [31-13](#)
    - configuring PIMv2 BSR [31-17](#)
    - monitoring mapping information [31-22](#)
    - using Auto-RP and BSR [31-21](#)
  - statistics, displaying system and network [31-50](#)
  - See also CGMP
  - See also DVMRP

## IP multicast routing (continued)

See also IGMP

See also PIM

## IP phones

and QoS [12-1](#)

automatic classification and queueing [26-29](#)

configuring [12-4](#)

ensuring port security with QoS [26-46](#)

trusted boundary for QoS [26-46](#)

## IP precedence [26-2](#)

## IP-precedence-to-DSCP map for QoS [26-60](#)

## IP protocols in ACLs [25-10](#)

## IP routes, monitoring [28-100](#)

## IP routing

connecting interfaces with [9-6](#)

enabling [28-17](#)

## IP traceroute

executing [34-12](#)

overview [34-11](#)

## IP unicast routing

address resolution [28-7](#)

administrative distances [28-89, 28-98](#)

ARP [28-8](#)

assigning IP addresses to Layer 3 interfaces [28-5](#)

authentication keys [28-99](#)

### broadcast

address [28-15](#)

flooding [28-15](#)

packets [28-12](#)

storms [28-12](#)

classless routing [28-6](#)

configuring static routes [28-88](#)

### default

addressing configuration [28-4](#)

gateways [28-10](#)

networks [28-89](#)

routes [28-89](#)

routing [28-2](#)

## IP unicast routing (continued)

directed broadcasts [28-13](#)

dynamic routing protocols [28-2](#)

enabling [28-17](#)

EtherChannel Layer 3 interface [28-3](#)

IGP [28-28](#)

inter-VLAN [28-2](#)

### IP addressing

classes [28-5](#)

configuring [28-4](#)

IRDP [28-11](#)

Layer 3 interfaces [28-3](#)

MAC address and IP address [28-7](#)

passive interfaces [28-97](#)

### protocols

distance-vector [28-2](#)

dynamic [28-2](#)

link-state [28-2](#)

proxy ARP [28-8](#)

redistribution [28-90](#)

reverse address resolution [28-7](#)

routed ports [28-3](#)

static routing [28-2](#)

steps to configure [28-3](#)

subnet mask [28-5](#)

subnet zero [28-6](#)

supernet [28-6](#)

UDP [28-14](#)

with SVIs [28-3](#)

See also BGP

See also EIGRP

See also IGRP

See also IS-IS

See also OSPF

See also RIP

## IRDP

configuring [28-11](#)

definition [28-11](#)

## IS-IS

- addresses [28-63](#)
- area routing [28-63](#)
- default configuration [28-67](#)
- monitoring [28-75](#)
- show commands [28-75](#)
- system routing [28-63](#)

## ISL

- and trunk ports [9-3](#)
- encapsulation [10-16](#)
- trunking with 802.1 tunneling [13-5](#)

## ISO CLNS

- clear commands [28-75](#)
- dynamic routing protocols [28-63](#)
- monitoring [28-75](#)
- NETs [28-63](#)
- NSAPs [28-63](#)
- OSI standard [28-63](#)

## ISO IGRP

- area routing [28-63](#)
- configuring [28-64](#)
- default configuration [28-64](#)
- interdomain routing [28-63](#)
- optional parameters [28-66](#)
- routing processes [28-64](#)
- system routing [28-63](#)

## J

- join messages, IGMP [17-2](#)

## K

## Kerberos

- authenticating to
  - boundary switch [7-34](#)
  - KDC [7-34](#)
  - network services [7-35](#)
- configuration examples [7-32](#)
- configuring [7-35](#)
- credentials [7-32](#)
- cryptographic software image [7-31](#)
- described [7-32](#)
- KDC [7-32](#)
- operation [7-34](#)
- realm [7-33](#)
- server [7-33](#)
- switch as trusted third party [7-32](#)
- terms [7-33](#)
- TGT [7-34](#)
- tickets [7-32](#)

## key distribution center

- See Kerberos, KDC

## L

- l2protocol-tunnel command [13-14](#)
- label binding [30-2](#)
- label distribution protocol
  - See LDP
- labels, MPLS [30-2](#)
- label switching router
  - See LSR
- LACP
  - See EtherChannel
- Layer 2 frames, classification with CoS [26-2](#)
- Layer 2 interfaces, default configuration [9-11](#)



- Layer 2 protocol tunneling
    - bypass mode [13-11, 13-13](#)
    - configuring [13-12](#)
    - default configuration [13-13](#)
    - defined [13-11](#)
    - guidelines [13-13](#)
  - Layer 2 traceroute
    - and ARP [34-11](#)
    - and CDP [34-10](#)
    - described [34-10](#)
    - IP addresses and subnets [34-11](#)
    - MAC addresses and VLANs [34-10](#)
    - multicast traffic [34-10](#)
    - multiple devices on a port [34-11](#)
    - unicast traffic [34-10](#)
    - usage guidelines [34-10](#)
  - Layer 2 trunks [10-17](#)
  - Layer 3 features [1-7](#)
  - Layer 3 interfaces
    - assigning IP addresses to [28-5](#)
    - changing from Layer 2 mode [28-5](#)
    - types of [28-3](#)
  - Layer 3 packets, classification methods [26-2](#)
  - LDAP [4-2](#)
  - LDP [30-6](#)
  - leave processing, IGMP [17-10](#)
  - lightweight directory access protocol
    - See LDAP
  - line configuration mode [2-3](#)
  - Link Aggregation Control Protocol
    - See EtherChannel
    - See LACP
  - links, unidirectional [20-1](#)
  - link state advertisements (LSAs) [28-32](#)
  - link-state protocols [28-2](#)
  - LLQ
    - described [26-29](#)
    - enabling [26-97](#)
  - load balancing, IGRP [28-24](#)
  - logging messages, ACL [25-9](#)
  - login authentication
    - with RADIUS [7-23](#)
    - with TACACS+ [7-14](#)
  - login banners [5-18](#)
  - log messages
    - See system message logging
  - Long-Reach Ethernet (LRE) technology [1-11](#)
  - loop guard
    - described [16-8](#)
    - enabling [16-15](#)
  - low-latency queueing
    - See LLQ
  - LSR [30-2](#)
- 
- ## M
- MAC addresses
    - aging time [5-22](#)
    - and VLAN association [5-21](#)
    - building the address table [5-21](#)
    - default configuration [5-21](#)
    - discovering [5-26](#)
    - displaying [5-25](#)
    - dynamic
      - learning [5-21](#)
      - removing [5-22](#)
    - in ACLs [25-25](#)
    - IP address association [28-7](#)
    - static
      - adding [5-25](#)
      - characteristics of [5-24](#)
      - removing [5-25](#)
  - MAC address-to-VLAN mapping [10-28](#)

- MAC extended access lists
  - applying to Layer 2 interfaces [25-26](#)
  - configuring for QoS [26-51](#)
  - creating [25-25](#)
  - defined [25-25](#)
  - for QoS classification [26-6](#)
- manageability features [1-3](#)
- management options
  - CLI [2-1](#)
  - CNS [4-1](#)
  - SNMP [24-1](#)
- mapping tables for QoS
  - configuring
    - CoS-to-DSCP [26-59](#)
    - DSCP [26-59](#)
    - DSCP-to-CoS [26-62](#)
    - DSCP-to-DSCP-mutation [26-63](#)
    - IP-precedence-to-DSCP [26-60](#)
    - policed-DSCP [26-61](#)
  - described [26-11](#)
- marking
  - action in
    - egress policy map [26-81, 26-84](#)
    - ingress policy map [26-54, 26-57](#)
  - described
    - egress [26-5, 26-24](#)
    - ingress [26-5, 26-10](#)
- matching, ACLs [25-6](#)
- maximum aging time
  - MSTP [15-21](#)
  - STP [14-21](#)
- maximum hop count, MSTP [15-21](#)
- maximum-paths command [28-50, 28-87](#)
- membership mode, VLAN port [10-3](#)
- messages
  - logging ACL violations [25-14](#)
  - to users through banners [5-18](#)
- metrics, in BGP [28-51](#)
- metric translations, between routing protocols [28-93](#)
- metro tags [13-2](#)
- MIBs
  - accessing files with FTP [A-3](#)
  - location of files [A-3](#)
  - overview [24-1](#)
  - SNMP interaction with [24-4](#)
  - supported [A-1](#)
- mirroring traffic for analysis [21-1](#)
- mismatches, autonegotiation [34-7](#)
- module number [9-7](#)
- monitoring
  - 802.1Q tunneling [13-16](#)
  - access groups [25-38](#)
  - ACL configuration [25-38](#)
  - BGP [28-62](#)
  - cables for unidirectional links [20-1](#)
  - CDP [19-5](#)
  - CEF [28-87](#)
  - EIGRP [28-42](#)
  - EoMPLS [30-22](#)
  - fallback bridging [33-10](#)
  - features [1-8](#)
  - HSRP [29-10](#)
  - IGMP
    - filters [17-22](#)
    - snooping [17-10](#)
  - interfaces [9-20](#)
  - IP
    - address tables [28-16](#)
    - multicast routing [31-49](#)
    - routes [28-100](#)
  - IS-IS [28-75](#)
  - ISO CLNS [28-75](#)
  - Layer 2 protocol tunneling [13-16](#)
  - MPLS [30-22](#)
  - MSDP peers [32-19](#)
  - multicast router interfaces [17-11](#)
  - multi-VRF CE [28-86](#)
  - MVR [17-18](#)

## monitoring (continued)

network traffic for analysis with probe [21-2](#)OSPF [28-36](#)

## port

blocking [18-14](#)protection [18-14](#)RP mapping information [31-22](#)source-active messages [32-19](#)speed and duplex mode [9-14](#)traffic flowing among switches [22-1](#)traffic suppression [18-14](#)tunneling [13-16](#)

## VLAN

filters [25-38](#)maps [25-38](#)VLANs [10-16](#)VMPS [10-33](#)VTP [11-16](#)MP-BGP [30-12](#)

## MPLS

configuring [30-7](#)default configuration [30-6](#)experimental field [30-19](#)label [30-2](#)monitoring [30-22](#)

## QoS

configuring [30-20](#)default configuration [30-20](#)experimental bits [30-19](#)uses [30-18](#)

## VPN

labels [30-3](#)packet flow [30-11](#)

## MSDP

benefits of [32-3](#)clearing MSDP connections and statistics [32-19](#)

## MSDP (continued)

controlling source information

forwarded by switch [32-12](#)originated by switch [32-8](#)received by switch [32-14](#)default configuration [32-4](#)

dense-mode regions

sending SA messages to [32-17](#)specifying the originating address [32-18](#)

filtering

incoming SA messages [32-14](#)SA messages to a peer [32-12](#)SA requests from a peer [32-11](#)join latency, defined [32-6](#)

meshed groups

configuring [32-16](#)defined [32-16](#)originating address, changing [32-18](#)overview [32-1](#)peer-RPF flooding [32-2](#)

peers

configuring a default [32-4](#)monitoring [32-19](#)peering relationship, overview [32-1](#)requesting source information from [32-8](#)shutting down [32-16](#)

source-active messages

caching [32-6](#)clearing cache entries [32-19](#)defined [32-2](#)filtering from a peer [32-11](#)filtering incoming [32-14](#)filtering to a peer [32-12](#)limiting data with TTL [32-14](#)monitoring [32-19](#)restricting advertised sources [32-9](#)

## MSTP

- boundary ports
  - configuration guidelines [15-13](#)
  - described [15-5](#)
- BPDU filtering
  - described [16-3](#)
  - enabling [16-12](#)
- BPDU guard
  - described [16-3](#)
  - enabling [16-11](#)
- CIST, described [15-3](#)
- configuration guidelines [15-12, 16-9](#)
- configuring
  - forward-delay time [15-20](#)
  - hello time [15-19](#)
  - link type for rapid convergence [15-22](#)
  - maximum aging time [15-21](#)
  - maximum hop count [15-21](#)
  - MST region [15-13](#)
  - path cost [15-18](#)
  - port priority [15-17](#)
  - root switch [15-14](#)
  - secondary root switch [15-16](#)
  - switch priority [15-19](#)
- CST
  - defined [15-3](#)
  - operations between regions [15-4](#)
- default configuration [15-12](#)
- default optional feature configuration [16-9](#)
- displaying status [15-23](#)
- enabling the mode [15-13](#)
- extended system ID
  - effects on root switch [15-14](#)
  - effects on secondary root switch [15-16](#)
  - unexpected behavior [15-15](#)
- instances supported [14-10](#)
- interface state, blocking to forwarding [16-2](#)
- interoperability and compatibility among modes [14-10](#)

## MSTP (continued)

- interoperability with 802.1D
  - described [15-5](#)
  - restarting migration process [15-22](#)
- IST
  - defined [15-3](#)
  - master [15-3](#)
  - operations within a region [15-3](#)
- loop guard
  - described [16-8](#)
  - enabling [16-15](#)
- mapping VLANs to MST instance [15-13](#)
- MST region
  - CIST [15-3](#)
  - configuring [15-13](#)
  - described [15-2](#)
  - hop-count mechanism [15-5](#)
  - IST [15-3](#)
  - supported spanning-tree instances [15-2](#)
- overview [15-2](#)
- Port Fast
  - described [16-2](#)
  - enabling [16-10](#)
- preventing root switch selection [16-7](#)
- root guard
  - described [16-7](#)
  - enabling [16-14](#)
- root switch
  - configuring [15-15](#)
  - effects of extended system ID [15-14](#)
  - unexpected behavior [15-15](#)
- shutdown Port Fast-enabled port [16-3](#)
- status, displaying [15-23](#)
- multicast groups
  - Immediate Leave [17-5](#)
  - joining [17-2](#)
  - leaving [17-4](#)
  - static joins [17-9](#)

multicast packets

- ACLs on [25-37](#)
- blocking [18-6](#)

multicast router interfaces, monitoring [17-11](#)

multicast router ports, adding [17-8](#)

Multicast Source Discovery Protocol

- See MSDP

multicast storm-control command [18-3](#)

multicast storms [18-1](#)

Multicast VLAN Registration

- See MVR

Multiple Spanning Tree Protocol

- See MSTP

multiple VPN routing/forwarding, customer edge devices

- See multi-VRF CE

multiple VPN routing/forwarding in customer edge devices

- See multi-VRF CE

multiprotocol label switching

- See MPLS

multi-VRF CE

- configuration example [28-82](#)
- configuration guidelines [28-79](#)
- configuring [28-78](#)
- default configuration [28-78](#)
- defined [28-76](#)
- displaying [28-86](#)
- monitoring [28-86](#)
- network components [28-78](#)
- packet-forwarding process [28-78](#)

MVR

- and address aliasing [17-15](#)
- configuring interfaces [17-17](#)
- default configuration [17-14](#)
- described [17-12](#)
- modes [17-16](#)
- monitoring [17-18](#)
- setting global parameters [17-15](#)

## N

named IP ACLs [25-13](#)

NameSpace Mapper

- See NSM

native VLAN

- and 802.1Q tunneling [13-4](#)
- configuring [10-24](#)
- default [10-24](#)

neighbor discovery/recovery, EIGRP [28-37](#)

neighbors, BGP [28-57](#)

network management

- CDP [19-1](#)
- RMON [22-1](#)
- SNMP [24-1](#)

Network Time Protocol

- See NTP

no commands [2-4](#)

non-IP traffic filtering [25-25](#)

nontrunking mode [10-18](#)

normal-range VLANs

- configuration modes [10-7](#)
- defined [10-1](#)

no switchport command [9-4](#)

note, described [xxxiv](#)

not-so-stubby areas

- See NSSA

NSAPs, as ISO IGRP addresses [28-63](#)

NSM [4-3](#)

NSSA, OSPF [28-32](#)

NTP

- and ACL time ranges [25-15](#)
- associations
  - authenticating [5-5](#)
  - defined [5-2](#)
  - enabling broadcast messages [5-7](#)
  - peer [5-6](#)
  - server [5-6](#)
- default configuration [5-4](#)

## NTP (continued)

- displaying the configuration [5-11](#)
- overview [5-2](#)
- restricting access
  - creating an access group [5-9](#)
  - disabling NTP services per interface [5-10](#)
- source IP address, configuring [5-10](#)
- stratum [5-2](#)
- synchronizing devices [5-6](#)
- time
  - services [5-2](#)
  - synchronizing [5-2](#)

NTP access group keywords, scanning order [5-9](#)

---

**O**

## Open Shortest Path First

See OSPF

optimizing system resources [6-1](#)

options, management [1-2](#)

## OSPF

- area parameters, configuring [28-32](#)
- configuring [28-30](#)
- default configuration
  - metrics [28-33](#)
  - route [28-33](#)
  - settings [28-29](#)
- described [28-28](#)
- interface parameters, configuring [28-31](#)
- LSA group pacing [28-35](#)
- monitoring [28-36](#)
- router IDs [28-35](#)
- route summarization [28-33](#)
- virtual links [28-33](#)

---

**P**

packet modification, with QoS [26-41](#)

## PAGP

See EtherChannel

parallel paths, in routing tables [28-87](#)

## passive interfaces

- configuring [28-97](#)
- OSPF [28-33](#)

## passwords

- default configuration [7-2](#)
- disabling recovery of [7-5](#)
- encrypting [7-3](#)
- overview [7-1](#)
- recovery of [34-3](#)

## setting

- enable [7-3](#)
- enable secret [7-3](#)
- Telnet [7-6](#)
- with usernames [7-6](#)

VTP domain [11-8](#)

## path cost

- MSTP [15-18](#)
- STP [14-18](#)

## PBR

- defined [28-94](#)
- enabling [28-95](#)
- fast-switched policy-based routing [28-96](#)
- local policy-based routing [28-96](#)

peers, BGP [28-57](#)

performance features [1-2](#)

## per-VLAN spanning-tree plus

See PVST+

PE to CE routing, configuring [28-81](#)

physical ports [9-2](#)

## PIM

- default configuration [31-8](#)
- dense mode
  - overview [31-4](#)
  - rendezvous point (RP), described [31-4](#)
  - RPF lookups [31-7](#)
- displaying neighbors [31-50](#)

## PIM (continued)

- enabling a mode [31-11](#)
- overview [31-3](#)
- router-query message interval, modifying [31-25](#)
- shared tree and source tree, overview [31-22](#)
- shortest path tree, delaying the use of [31-24](#)
- sparse mode
  - join messages and shared tree [31-4](#)
  - overview [31-4](#)
  - prune messages [31-5](#)
  - RPF lookups [31-7](#)
- versions
  - interoperability [31-9](#)
  - troubleshooting interoperability problems [31-22](#)
  - v2 improvements [31-4](#)

PIM-DVMRP, as snooping method [17-7](#)

## ping

- character output description [34-9](#)
- executing [34-9](#)
- overview [34-8](#)

poison-reverse updates, IGRP [28-24](#)policed-DSCP map for QoS [26-61](#)

## policers

- configuring
  - egress, two-rate [26-80](#)
  - for each matched traffic class [26-54](#)
  - for more than one traffic class [26-57](#)
- displaying aggregate [26-75](#)
- egress, two-rate [26-24](#)
- ingress, single-rate [26-10](#)
- number supported [26-40, 26-77](#)
- types of ingress [26-9](#)

## policing

- egress, described [26-24](#)
- ingress, described [26-9](#)
- token-bucket algorithm [26-10, 26-24](#)

## policy-based routing

- See PBR

## policy maps for QoS

- characteristics of ingress [26-54](#)
- described [26-8, 26-23](#)
- displaying [26-75, 26-101](#)

## port ACLs

- defined [25-2](#)
- types of [25-3](#)

## Port Aggregation Protocol

- See EtherChannel

## port-based authentication

- authentication server
  - defined [8-2](#)
  - RADIUS server [8-2](#)
- client, defined [8-2](#)
- configuration guidelines [8-11](#)
- configuring
  - 802.1x authentication [8-11](#)
  - guest VLAN [8-18](#)
  - host mode [8-17](#)
  - manual re-authentication of a client [8-14](#)
  - periodic re-authentication [8-14](#)
  - quiet period [8-15](#)
  - RADIUS server [8-14](#)
  - RADIUS server parameters on the switch [8-13](#)
  - switch-to-client frame-retransmission number [8-16](#)
  - switch-to-client retransmission time [8-15](#)
- default configuration [8-10](#)
- described [8-1](#)
- device roles [8-2](#)
- displaying statistics [8-19](#)
- EAPOL-start frame [8-3](#)
- EAP-request/identity frame [8-3](#)
- EAP-response/identity frame [8-3](#)
- encapsulation [8-2](#)
- guest VLAN
  - configuration guidelines [8-8](#)
  - described [8-8](#)
- initiation and message exchange [8-3](#)
- method lists [8-11](#)

- port-based authentication (continued)
  - multiple-hosts mode, described [8-17](#)
- per-user ACLs
  - AAA authorization [8-11](#)
  - configuration tasks [8-9](#)
  - described [8-8](#)
  - RADIUS server attributes [8-9](#)
- ports
  - authorization state and dot1x port-control command [8-4](#)
  - authorized and unauthorized [8-4](#)
  - voice VLAN [8-6](#)
- port security
  - and voice VLAN [8-6](#)
  - described [8-6](#)
  - interactions [8-6](#)
  - multiple-hosts mode [8-17](#)
- resetting to default values [8-18](#)
- statistics, displaying [8-19](#)
- switch
  - as proxy [8-2](#)
  - RADIUS client [8-2](#)
- topologies, supported [8-5](#)
- VLAN assignment
  - AAA authorization [8-11](#)
  - characteristics [8-7](#)
  - configuration tasks [8-8](#)
  - described [8-7](#)
- voice VLAN
  - described [8-6](#)
  - PVID [8-6](#)
  - VVID [8-6](#)
- port blocking [18-5](#)
- port-channel
  - See EtherChannel
- Port Fast
  - described [16-2](#)
  - enabling [16-10](#)
  - mode, spanning tree [10-29](#)
- port membership modes, VLAN [10-3](#)
- port priority
  - MSTP [15-17](#)
  - STP [14-17](#)
- ports
  - 802.1Q tunnel [10-4](#)
  - access [9-2](#)
  - blocking [18-5](#)
  - configuring [9-6](#)
  - dynamic access [10-4](#)
  - enhanced services (ES) [9-7](#)
  - identifying [9-6](#)
  - numbering [9-7](#)
  - protected [18-4](#)
  - routed [9-4](#)
  - secure [18-6](#)
  - static-access [10-3, 10-11](#)
  - switch [9-2](#)
  - trunks [10-3, 10-16](#)
  - VLAN assignments [10-11](#)
- port security
  - aging [18-13](#)
  - and QoS trusted boundary [26-46](#)
  - configuration guidelines [18-9](#)
  - configuring [18-10](#)
  - default configuration [18-9](#)
  - described [18-6](#)
  - displaying [18-14](#)
  - on trunk ports [18-11](#)
  - sticky learning [18-7](#)
  - violations [18-8](#)
- port-shutdown response, VMPS [10-28](#)
- preferential treatment of traffic
  - See QoS
- prefix lists, BGP [28-54](#)
- preventing unauthorized access [7-1](#)



- priority
    - HSRP [29-6](#)
    - overriding CoS [12-5](#)
    - trusting CoS [12-5](#)
  - priority queues for QoS
    - egress LLQ [26-97](#)
    - ingress [26-68](#)
  - private VLAN edge ports
    - See protected ports
  - privileged EXEC mode [2-2](#)
  - privilege levels
    - changing the default for lines [7-9](#)
    - exiting [7-10](#)
    - logging into [7-10](#)
    - overview [7-2, 7-8](#)
    - setting a command with [7-8](#)
  - protected ports [18-4](#)
  - protocol-dependent modules, EIGRP [28-38](#)
  - Protocol-Independent Multicast Protocol
    - See PIM
  - provider edge devices
    - and MPLS [30-7](#)
    - and MPLS labels [30-3](#)
    - in VPNs [30-5](#)
    - using multi-VRF CE [28-76](#)
  - proxy ARP
    - configuring [28-10](#)
    - definition [28-8](#)
    - with IP routing disabled [28-10](#)
  - pruning, VTP
    - enabling [11-14](#)
    - enabling on a port [10-23](#)
    - examples [11-5](#)
    - overview [11-4](#)
  - pruning-eligible list
    - changing [10-23](#)
    - for VTP pruning [11-4](#)
    - VLANs [11-14](#)
  - publications for products, technologies, and network solutions [xxxviii](#)
  - PVST+
    - 802.1Q trunking interoperability [14-10](#)
    - described [14-9](#)
    - instances supported [14-10](#)
- 
- ## Q
- QoS
    - ACLs
      - IP extended [26-50](#)
      - IP standard [26-49](#)
      - MAC [26-51](#)
    - aggregate policers, configuring [26-57](#)
    - and 802.1Q tunneling [26-44, 26-77](#)
    - auto-QoS
      - categorizing traffic [26-30](#)
      - configuration and defaults display [26-37](#)
      - configuration guidelines [26-33](#)
      - described [26-29](#)
      - disabling [26-34](#)
      - effects on running configuration [26-33](#)
      - egress queue-set defaults [26-30](#)
      - enabling for VoIP [26-34](#)
      - example configuration [26-35](#)
      - generated commands, displaying [26-34](#)
      - ingress queue defaults [26-30](#)
      - initial configuration, displaying [26-37](#)
      - list of generated commands [26-31](#)
    - basic model [26-4](#)
    - classification
      - class maps, described [26-8, 26-23](#)
      - egress, defined [26-5](#)
      - flowchart, ingress [26-7](#)
      - forwarding treatment [26-3](#)
      - in frames and packets [26-3](#)
      - ingress, defined [26-4](#)
      - IP ACLs, described [26-6, 26-8](#)

## QoS (continued)

## classification (continued)

- MAC ACLs, described [26-6, 26-8](#)
- options for IP traffic [26-6](#)
- options for non-IP traffic [26-6](#)
- trust DSCP, described [26-6](#)
- trusted CoS, described [26-6](#)
- trust IP precedence, described [26-6](#)

## class maps

- configuring egress [26-78](#)
- configuring ingress [26-52](#)
- displaying [26-75, 26-101](#)

## configuration guidelines

- auto-QoS [26-33](#)
- hierarchical QoS [26-76](#)
- standard QoS [26-40](#)

## default configuration

- auto-QoS [26-30](#)
- hierarchical QoS [26-76](#)
- standard QoS [26-38](#)

displaying statistics [26-75](#)DSCP maps [26-59](#)

## egress hierarchical queues

- average queue size calculation [26-28](#)
- bandwidth limited stream [26-19](#)
- CBWFQ [26-28](#)
- congestion control [26-27](#)
- default queue [26-27](#)
- described [26-5, 26-27](#)
- LLQ [26-29](#)
- number of queues supported [26-27](#)
- queue creation [26-23](#)
- scheduling [26-5](#)
- See also QoS, hierarchical QoS
- tail drop [26-27](#)
- WRED [26-28](#)

## QoS (continued)

## egress queue-sets

- allocating buffer space [26-69](#)
- buffer allocation scheme, described [26-18](#)
- characteristics [26-69](#)
- configuring shaped weights for SRR [26-72](#)
- configuring shared weights for SRR [26-73](#)
- described [26-5](#)
- displaying the threshold map [26-72](#)
- flowchart [26-17](#)
- mapping DSCP or CoS values [26-71](#)
- scheduling, described [26-5](#)
- setting WTD thresholds [26-69](#)
- WTD, described [26-19](#)

enabling globally [26-42](#)

## flowcharts

- egress, two-rate policing and marking [26-24](#)
- egress hierarchical queues queueing and scheduling [26-26](#)
- egress queue-set queueing and scheduling [26-17](#)
- ingress, single-rate policing and marking [26-11](#)
- ingress classification [26-7](#)
- ingress queueing and scheduling [26-15](#)

## hierarchical levels

- class level, described [26-20](#)
- physical interface level, described [26-22](#)
- supported number of class-level classes [26-20](#)
- supported number of VLAN-level classes [26-21](#)
- VLAN level, described [26-21](#)

## hierarchical QoS

- CBWFQ [26-28](#)
- CBWFQ and DSCP-based WRED [26-89](#)
- CBWFQ and IP precedence-based WRED [26-93](#)
- CBWFQ and tail drop [26-86](#)
- child policy [26-21](#)
- classification based on class maps [26-23, 26-78](#)
- configuration guidelines [26-76](#)
- congestion avoidance [26-27](#)
- congestion management [26-27](#)

## QoS (continued)

## hierarchical QoS (continued)

- default class [26-23](#)
- default configuration [26-76](#)
- displaying [26-101](#)
- LLQ [26-29, 26-97](#)
- marking [26-24, 26-84](#)
- matching criteria [26-23](#)
- policing, described [26-24](#)
- See also QoS, egress hierarchical queues
- See also QoS, hierarchical levels
- shaping [26-29, 26-99](#)
- tail drop [26-27](#)
- traffic policies, described [26-23](#)
- two-rate traffic policer [26-24, 26-80](#)
- WRED [26-28](#)

implicit deny [26-8](#)

## ingress queues

- allocating bandwidth [26-67](#)
- allocating buffer space [26-66](#)
- buffer and bandwidth allocation, described [26-16](#)
- characteristics [26-64](#)
- configuring shared weights for SRR [26-67](#)
- configuring the priority queue [26-68](#)
- described [26-5](#)
- displaying the threshold map [26-65](#)
- flowchart [26-15](#)
- mapping DSCP or CoS values [26-65](#)
- priority queue, described [26-16](#)
- scheduling, described [26-5](#)
- setting WTD thresholds [26-65](#)
- WTD, described [26-16](#)

in MPLS networks [30-18](#)

## IP phones

- automatic classification and queueing [26-29](#)
- detection and trusted settings [26-29, 26-46](#)
- limiting bandwidth on egress interface [26-74](#)

## QoS (continued)

## mapping tables

- CoS-to-DSCP [26-59](#)
- displaying [26-75](#)
- DSCP-to-CoS [26-62](#)
- DSCP-to-DSCP-mutation [26-63](#)
- IP-precedence-to-DSCP [26-60](#)
- policed-DSCP [26-61](#)
- types of [26-11](#)

marked-down actions [26-55, 26-81, 26-84](#)

## marking

- described, egress [26-5, 26-24](#)
- described, ingress [26-5, 26-10](#)

overview [26-2](#)packet modification [26-41](#)

## policers

- configuring [26-55, 26-57](#)
- displaying aggregate [26-75](#)
- egress, two-rate [26-24](#)
- ingress, single-rate [26-10](#)
- number supported [26-40, 26-77](#)
- types of ingress [26-9](#)

policies, attaching to an interface [26-10, 26-25](#)

## policing

- described [26-4](#)
- egress, described [26-24](#)
- ingress, described [26-9](#)
- token-bucket algorithm [26-10](#)

## policy maps

- characteristics of ingress [26-54](#)
- configuring [26-54](#)
- described [26-8, 26-23](#)
- displaying [26-75, 26-101](#)

QoS label, defined [26-4](#)

## QoS (continued)

## queues

- configuring egress queue-set characteristics [26-69](#)

- configuring ingress characteristics [26-64](#)

- location of [26-12](#)

- See also QoS, egress hierarchical queues

- See also QoS, egress queue-sets

- See also QoS, ingress queues

- SRR, described [26-14](#)

- WTD, described [26-13](#)

- rewrites [26-41](#)

- trusted boundary, configuring [26-46](#)

## trust states

- bordering another domain [26-47](#)

- described [26-6](#)

- trusted device [26-46](#)

- within the domain [26-42](#)

- QoS features [1-6](#)

## quality of service

- See QoS

- queries, IGMP [17-3](#)

## R

## RADIUS

## attributes

- vendor-proprietary [7-30](#)

- vendor-specific [7-29](#)

## configuring

- accounting [7-28](#)

- authentication [7-23](#)

- authorization [7-27](#)

- communication, global [7-21, 7-28](#)

- communication, per-server [7-20, 7-21](#)

- multiple UDP ports [7-21](#)

- default configuration [7-20](#)

- defining AAA server groups [7-25](#)

- displaying the configuration [7-31](#)

- identifying the server [7-20](#)

## RADIUS (continued)

- limiting the services to the user [7-27](#)

- method list, defined [7-20](#)

- operation of [7-19](#)

- overview [7-18](#)

- suggested network environments [7-18](#)

- tracking services accessed by user [7-28](#)

## range

- macro [9-9](#)

- of interfaces [9-8](#)

- rapid convergence [15-7](#)

- rapid per-VLAN spanning-tree plus

- See rapid PVST+

- rapid PVST+

- 802.1Q trunking interoperability [14-10](#)

- described [14-9](#)

- instances supported [14-10](#)

- Rapid Spanning Tree Protocol

- See RSTP

- RARP [28-8](#)

## RCP

- configuration files

- downloading [B-16](#)

- overview [B-15](#)

- preparing the server [B-16](#)

- uploading [B-17](#)

- image files

- deleting old image [B-31](#)

- downloading [B-29](#)

- preparing the server [B-28](#)

- uploading [B-31](#)

- reconfirmation interval, VMPS, changing [10-32](#)

- recovery procedures [34-1](#)

- redundancy

- EtherChannel [27-2](#)

- HSRP [29-1](#)

## redundancy (continued)

### STP

backbone [14-8](#)

path cost [10-26](#)

port priority [10-25](#)

redundant links and UplinkFast [16-13](#)

reliable transport protocol, EIGRP [28-37](#)

reloading software [3-15](#)

Remote Authentication Dial-In User Service

See RADIUS

Remote Copy Protocol

See RCP

Remote Network Monitoring

See RMON

Remote SPAN

See RSPAN

resets, in BGP [28-49](#)

resetting a UDLD-shutdown interface [20-6](#)

restricting access

NTP services [5-8](#)

overview [7-1](#)

passwords and privilege levels [7-2](#)

RADIUS [7-17](#)

TACACS+ [7-10](#)

retry count, VMPS, changing [10-32](#)

reverse address resolution [28-7](#)

Reverse Address Resolution Protocol

See RARP

## RFC

1058, RIP [28-18](#)

1112, IP multicast and IGMP [17-2](#)

1157, SNMPv1 [24-2](#)

1163, BGP [28-43](#)

1166, IP addresses [28-5](#)

1253, OSPF [28-28](#)

1267, BGP [28-43](#)

1305, NTP [5-2](#)

1587, NSSAs [28-28](#)

1757, RMON [22-2](#)

## RFC (continued)

1771, BGP [28-43](#)

1901, SNMPv2C [24-2](#)

1902 to 1907, SNMPv2 [24-2](#)

2236, IP multicast and IGMP [17-2](#)

2273-2275, SNMPv3 [24-2](#)

## RIP

advertisements [28-18](#)

authentication [28-21](#)

configuring [28-19](#)

default configuration [28-18](#)

described [28-18](#)

hop counts [28-18](#)

split horizon [28-21](#)

summary addresses [28-21](#)

## RMON

default configuration [22-3](#)

displaying status [22-6](#)

enabling alarms and events [22-3](#)

groups supported [22-2](#)

overview [22-1](#)

statistics

collecting group Ethernet [22-6](#)

collecting group history [22-5](#)

root guard

described [16-7](#)

enabling [16-14](#)

root switch

MSTP [15-14](#)

STP [14-14](#)

route calculation timers, OSPF [28-34](#)

route dampening, BGP [28-61](#)

routed packets, ACLs on [25-36](#)

routed ports

configuring [28-3](#)

defined [9-4](#)

IP addresses on [9-18, 28-3](#)

route-map command [28-95](#)

- route maps
    - BGP [28-52](#)
    - policy-based routing [28-94](#)
  - router ACLs
    - defined [25-2](#)
    - types of [25-3](#)
  - route reflectors, BGP [28-60](#)
  - router ID, OSPF [28-35](#)
  - route selection, BGP [28-50](#)
  - route summarization, OSPF [28-33](#)
  - route targets, VPN [28-78](#)
  - routing
    - default [28-2](#)
    - dynamic [28-2](#)
    - redistribution of information [28-90](#)
    - static [28-2](#)
  - routing domain confederation, BGP [28-59](#)
  - Routing Information Protocol
    - See RIP
  - routing protocol administrative distances [28-89](#)
  - RSPAN
    - characteristics [21-8](#)
    - configuration guidelines [21-16](#)
    - default configuration [21-9](#)
    - destination ports [21-7](#)
    - displaying status [21-23](#)
    - interaction with other features [21-8](#)
    - monitored ports [21-5](#)
    - monitoring ports [21-7](#)
    - overview [21-1](#)
    - received traffic [21-4](#)
    - session limits [21-10](#)
    - sessions
      - creating [21-17](#)
      - defined [21-3](#)
      - limiting source traffic to specific VLANs [21-22](#)
      - specifying monitored ports [21-17](#)
      - with ingress traffic enabled [21-20](#)
  - RSPAN (continued)
    - source ports [21-5](#)
    - transmitted traffic [21-5](#)
    - VLAN-based [21-6](#)
  - RSTP
    - active topology, determining [15-6](#)
    - BPDU
      - format [15-9](#)
      - processing [15-10](#)
    - designated port, defined [15-6](#)
    - designated switch, defined [15-6](#)
    - interoperability with 802.1D
      - described [15-5](#)
      - restarting migration process [15-22](#)
      - topology changes [15-10](#)
    - overview [15-6](#)
    - port roles
      - described [15-6](#)
      - synchronized [15-8](#)
    - proposal-agreement handshake process [15-7](#)
    - rapid convergence
      - described [15-7](#)
      - edge ports and Port Fast [15-7](#)
      - point-to-point links [15-7, 15-22](#)
      - root ports [15-7](#)
    - root port, defined [15-6](#)
    - See also MSTP
  - running configuration, saving [3-10](#)
- 
- ## S
- scheduled reloads [3-15](#)
  - scheduling
    - egress hierarchical queues
      - CBWFQ [26-28, 26-86](#)
      - LLQ [26-29, 26-97](#)
      - shaping [26-29, 26-99](#)

- scheduling (continued)
  - egress queue-sets
    - shaped or shared mode [26-19](#)
    - shaped weights [26-72](#)
    - shared weights [26-73](#)
    - WTD thresholds [26-19, 26-69](#)
  - ingress queues
    - priority queueing [26-16, 26-68](#)
    - WTD thresholds [26-16, 26-65](#)
- SDM
  - described [6-1](#)
  - templates
    - configuration guidelines [6-3](#)
    - configuring [6-2, 6-3](#)
    - number of [6-1](#)
- secure MAC addresses
  - deleting [18-12](#)
  - maximum number of [18-8](#)
  - types of [18-7](#)
- secure ports, configuring [18-6](#)
- secure remote connections [7-37](#)
- Secure Shell
  - See SSH
- security, port [18-6](#)
- security features [1-5](#)
- sequence numbers in log messages [23-7](#)
- server mode, VTP [11-3](#)
- service-provider network
  - and 802.1Q tunneling [13-2](#)
  - and customer VLANs [13-2](#)
  - and EoMPLS [30-12](#)
  - and MPLS [30-2](#)
  - configuring MPLS VPNs [30-6](#)
  - Layer 2 protocols across [13-11](#)
  - MSTP and RSTP [15-1](#)
  - VPNs in [30-3](#)
- set-request operation [24-4](#)
- setup (CLI) program
  - described [3-2](#)
  - See also hardware installation guide
- severity levels, defining in system messages [23-8](#)
- SFPs
  - interface numbering [9-6](#)
  - security and identification [34-8](#)
- shaped round robin
  - See SRR
- shaping, average-rate
  - configuring [26-99](#)
  - described [26-29](#)
- show access-lists hw-summary command [25-19](#)
- show and more command output, filtering [2-9](#)
- show cdp traffic command [19-5](#)
- show configuration command [9-17](#)
- show forward command [34-14](#)
- show interfaces command [9-14, 9-17](#)
- show l2protocol command [13-15](#)
- show platform forward command [34-14](#)
- show running-config command
  - displaying ACLs [25-18, 25-38](#)
  - interface description in [9-17](#)
- shutdown command on interfaces [9-22](#)
- shutdown threshold for Layer 2 protocol packets [13-13](#)
- Simple Network Management Protocol
  - See SNMP
- SNAP [19-1](#)
- SNMP
  - accessing MIB variables with [24-4](#)
  - agent
    - described [24-4](#)
    - disabling [24-7](#)
  - authentication level [24-9](#)
  - community strings
    - configuring [24-7](#)
    - overview [24-4](#)
  - configuration examples [24-14](#)
  - configuration guidelines [24-6](#)

## SNMP (continued)

- default configuration [24-6](#)
  - engine ID [24-6](#)
  - groups [24-6, 24-8](#)
  - host [24-6](#)
  - informs
    - and trap keyword [24-10](#)
    - described [24-5](#)
    - differences from traps [24-5](#)
    - enabling [24-13](#)
  - limiting access by TFTP servers [24-13](#)
  - limiting system log messages to NMS [23-9](#)
  - manager functions [24-3](#)
  - MIBs
    - location of [A-3](#)
    - supported [A-1](#)
  - notifications [24-5](#)
  - overview [24-1, 24-4](#)
  - status, displaying [24-15](#)
  - system contact and location [24-13](#)
  - trap manager, configuring [24-12](#)
  - traps
    - described [24-3, 24-5](#)
    - differences from informs [24-5](#)
    - enabling [24-10](#)
    - enabling MAC address notification [5-23](#)
    - overview [24-1, 24-4](#)
    - types of [24-10](#)
  - users [24-6, 24-8](#)
  - versions supported [24-2](#)
- SNMPv1 [24-2](#)
- SNMPv2C [24-2](#)
- SNMPv3 [24-2](#)
- snooping, IGMP [17-2](#)

## software images

- location in flash memory [B-20](#)
  - recovery procedures [34-2](#)
  - scheduling reloads [3-15](#)
  - tar file format, described [B-20](#)
  - See also downloading and uploading
  - source addresses, in ACLs [25-10](#)
  - source-and-destination-IP address based forwarding, EtherChannel [27-7](#)
  - source-and-destination MAC address forwarding, EtherChannel [27-7](#)
  - source-IP address based forwarding, EtherChannel [27-7](#)
  - source-MAC address forwarding, EtherChannel [27-7](#)
- SPAN
- configuration guidelines [21-10](#)
  - default configuration [21-9](#)
  - destination ports [21-7](#)
  - displaying status [21-23](#)
  - interaction with other features [21-8](#)
  - monitored ports [21-5](#)
  - monitoring ports [21-7](#)
  - overview [21-1](#)
  - received traffic [21-4](#)
  - session limits [21-10](#)
  - sessions
    - configuring ingress forwarding [21-14, 21-21](#)
    - creating [21-11](#)
    - defined [21-3](#)
    - limiting source traffic to specific VLANs [21-15](#)
    - removing destination (monitoring) ports [21-12](#)
    - specifying monitored ports [21-11](#)
    - with ingress traffic enabled [21-13](#)
  - source ports [21-5](#)
  - transmitted traffic [21-5](#)
  - VLAN-based [21-6](#)



- spanning tree and native VLANs [10-19](#)
- Spanning Tree Protocol
  - See STP
- SPAN traffic [21-4](#)
- speed, configuring on interfaces [9-12](#)
- split horizon
  - IGRP [28-27](#)
  - RIP [28-21](#)
- SRR
  - configuring
    - shaped weights on egress queue-sets [26-72](#)
    - shared weights on egress queue-sets [26-73](#)
    - shared weights on ingress queues [26-67](#)
  - described [26-14](#)
  - shaped mode [26-14](#)
  - shared mode [26-14](#)
- SSH
  - configuring [7-37](#)
  - cryptographic software image [7-37](#)
  - described [7-37](#)
  - displaying settings [7-37](#)
- standby ip command [29-5](#)
- standby router [29-1](#)
- standby timers, HSRP [29-8](#)
- startup configuration
  - booting
    - manually [3-12](#)
    - specific image [3-12](#)
  - clearing [B-18](#)
  - configuration file
    - automatically downloading [3-11](#)
    - specifying the filename [3-11](#)
  - default boot configuration [3-11](#)
- static access ports
  - assigning to VLAN [10-11](#)
  - defined [9-3, 10-3](#)
- static addresses
  - See addresses
- static routes, configuring [28-88](#)
- static routing [28-2](#)
- static VLAN membership [10-2](#)
- statistics
  - 802.1x [8-19](#)
  - CDP [19-5](#)
  - interface [9-21](#)
  - IP multicast routing [31-50](#)
  - OSPF [28-36](#)
  - QoS ingress and egress [26-75](#)
  - RMON group Ethernet [22-6](#)
  - RMON group history [22-5](#)
  - SNMP input and output [24-15](#)
  - VTP [11-16](#)
- sticky secure MAC address learning [18-7](#)
- storm control
  - configuring [18-3](#)
  - described [18-1](#)
  - displaying [18-14](#)
  - thresholds [18-1](#)
- STP
  - 802.1D and bridge ID [14-4](#)
  - 802.1D and multicast addresses [14-8](#)
  - 802.1t and VLAN identifier [14-4](#)
  - accelerating root port selection [16-4](#)
  - BackboneFast
    - described [16-5](#)
    - enabling [16-14](#)
  - BPDU filtering
    - described [16-3](#)
    - enabling [16-12](#)
  - BPDU guard
    - described [16-3](#)
    - enabling [16-11](#)
  - BPDU message exchange [14-3](#)
  - configuration guidelines [14-12, 16-9](#)

## STP (continued)

## configuring

- forward-delay time [14-21](#)
- hello time [14-20](#)
- maximum aging time [14-21](#)
- path cost [14-18](#)
- port priority [14-17](#)
- root switch [14-14](#)
- secondary root switch [14-16](#)
- spanning-tree mode [14-13](#)
- switch priority [14-19](#)

counters, clearing [14-22](#)default configuration [14-11](#)default optional feature configuration [16-9](#)designated port, defined [14-3](#)designated switch, defined [14-3](#)detecting indirect link failures [16-6](#)disabling [14-14](#)displaying status [14-22](#)

## extended system ID

- effects on root switch [14-14](#)
- effects on the secondary root switch [14-16](#)
- overview [14-4](#)
- unexpected behavior [14-15](#)

inferior BPDU [14-3](#)instances supported [14-10](#)interface state, blocking to forwarding [16-2](#)

## interface states

- blocking [14-5](#)
- disabled [14-7](#)
- forwarding [14-5, 14-6](#)
- learning [14-6](#)
- listening [14-6](#)
- overview [14-4](#)

interoperability and compatibility among modes [14-10](#)Layer 2 protocol tunneling [13-10](#)limitations with 802.1Q trunks [14-10](#)

## STP (continued)

## load sharing

- overview [10-24](#)
- using path costs [10-26](#)
- using port priorities [10-25](#)

## loop guard

- described [16-8](#)
- enabling [16-15](#)

modes supported [14-9](#)multicast addresses, effect of [14-8](#)overview [14-2](#)path costs [10-26, 10-27](#)

## Port Fast

- described [16-2](#)
- enabling [16-10](#)

port priorities [10-25](#)preventing root switch selection [16-7](#)protocols supported [14-9](#)redundant connectivity [14-8](#)

## root guard

- described [16-7](#)
- enabling [16-14](#)

root port, defined [14-3](#)

## root switch

- configuring [14-15](#)
- effects of extended system ID [14-4, 14-14](#)
- election [14-3](#)
- unexpected behavior [14-15](#)

shutdown Port Fast-enabled port [16-3](#)status, displaying [14-22](#)superior BPDU [14-3](#)timers, described [14-20](#)

## UplinkFast

- described [16-4](#)
- enabling [16-13](#)

VLAN-bridge [14-11](#)

- stratum, NTP [5-2](#)
- stub areas, OSPF [28-32](#)
- subnet mask [28-5](#)
- subnet zero [28-6](#)
- success response, VMPS [10-28](#)
- summer time [5-13](#)
- SunNet Manager [1-2](#)
- supernet [28-6](#)
- SVIs
  - and IP unicast routing [28-3](#)
  - and router ACLs [25-3](#)
  - connecting VLANs [9-5](#)
  - defined [9-4](#)
  - routing between VLANs [10-2](#)
- switch console port [1-3](#)
- Switch Database Management
  - See SDM
- switched packets, ACLs on [25-35](#)
- Switched Port Analyzer
  - See SPAN
- switched ports [9-2](#)
- switchport block multicast command [18-6](#)
- switchport block unicast command [18-6](#)
- switchport command [9-11](#)
- switchport mode dot1q-tunnel command [13-6](#)
- switchport protected command [18-5](#)
- switch priority
  - MSTP [15-19](#)
  - STP [14-19](#)
- switch software features [1-1](#)
- switch virtual interface
  - See SVI
- synchronization, BGP [28-47](#)
- syslog
  - See system message logging
- system clock
  - configuring
    - daylight saving time [5-13](#)
    - manually [5-11](#)
    - summer time [5-13](#)
    - time zones [5-12](#)
  - displaying the time and date [5-12](#)
  - overview [5-2](#)
  - See also NTP
- system message logging
  - default configuration [23-3](#)
  - defining error message severity levels [23-8](#)
  - disabling [23-4](#)
  - displaying the configuration [23-12](#)
  - enabling [23-4](#)
  - facility keywords, described [23-11](#)
  - level keywords, described [23-8](#)
  - limiting messages [23-9](#)
  - message format [23-2](#)
  - overview [23-1](#)
  - sequence numbers, enabling and disabling [23-7](#)
  - setting the display destination device [23-4](#)
  - synchronizing log messages [23-5](#)
  - timestamps, enabling and disabling [23-7](#)
  - UNIX syslog servers
    - configuring the daemon [23-10](#)
    - configuring the logging facility [23-11](#)
    - facilities supported [23-11](#)
- system MTU
  - and 802.1Q tunneling [13-5](#)
  - and EoMPLS [30-14](#)
  - and IS-IS LSPs [28-70](#)
  - configuring [9-19](#)
  - maximum size supported [9-19](#)
- system name
  - default configuration [5-15](#)
  - default setting [5-15](#)
  - manual configuration [5-15](#)
  - See also DNS

system prompt  
     default setting [5-15](#)  
     manual configuration [5-16](#)  
 system resources, optimizing [6-1](#)  
 system routes, IGRP [28-23](#)  
 system routing  
     IS-IS [28-63](#)  
     ISO IGRP [28-63](#)

---

## T

TACACS+  
     accounting, defined [7-11](#)  
     authentication, defined [7-11](#)  
     authorization, defined [7-11](#)  
     configuring  
         accounting [7-16](#)  
         authentication key [7-13](#)  
         authorization [7-16](#)  
         login authentication [7-14](#)  
     default configuration [7-13](#)  
     displaying the configuration [7-17](#)  
     identifying the server [7-13](#)  
     limiting the services to the user [7-16](#)  
     operation of [7-12](#)  
     overview [7-10](#)  
     tracking services accessed by user [7-16](#)  
 TAC website [xxxvii](#)  
 tag distribution protocol  
     See TDP  
 tagged packets  
     802.1Q [13-3](#)  
     Layer 2 protocol [13-10](#)  
 tail drop  
     configuring [26-86](#)  
     described [26-27](#)

tar files  
     creating [B-6](#)  
     displaying the contents of [B-6](#)  
     extracting [B-7](#)  
     image file format [B-20](#)  
 TDP [30-6](#)  
 technical assistance  
     case priority definitions [xxxvii](#)  
     opening a case [xxxvii](#)  
     TAC website [xxxvi](#)  
 Telnet  
     accessing management interfaces [2-10](#)  
     from a browser [2-10](#)  
     number of connections [1-3](#)  
     setting a password [7-6](#)  
 templates, SDM [6-2](#)  
 Terminal Access Controller Access Control System Plus  
     See TACACS+  
 terminal lines, setting a password [7-6](#)  
 TFTP  
     configuration files  
         downloading [B-11](#)  
         preparing the server [B-10](#)  
         uploading [B-11](#)  
     configuration files in base directory [3-6](#)  
     configuring for autoconfiguration [3-5](#)  
     image files  
         deleting [B-23](#)  
         downloading [B-22](#)  
         preparing the server [B-21](#)  
         uploading [B-23](#)  
     limiting access by servers [24-13](#)  
 threshold, traffic level [18-2](#)  
 time  
     See NTP and system clock  
 time-range command [25-15](#)  
 time ranges in ACLs [25-15](#)  
 timestamps in log messages [23-7](#)  
 time zones [5-12](#)

- Token Ring VLANs
  - support for [10-6](#)
  - VTP support [11-4](#)
- traceroute, Layer 2
  - and ARP [34-11](#)
  - and CDP [34-10](#)
  - described [34-10](#)
  - IP addresses and subnets [34-11](#)
  - MAC addresses and VLANs [34-10](#)
  - multicast traffic [34-10](#)
  - multiple devices on a port [34-11](#)
  - unicast traffic [34-10](#)
  - usage guidelines [34-10](#)
- traceroute command [34-12](#)
  - See also IP traceroute
- traffic
  - blocking flooded [18-6](#)
  - fragmented [25-5](#)
  - unfragmented [25-5](#)
- traffic shaping
  - See shaping, average-rate
- traffic suppression [18-1](#)
- transparent mode, VTP [11-3, 11-12](#)
- trap-door mechanism [3-2](#)
- traps
  - configuring MAC address notification [5-23](#)
  - configuring managers [24-10](#)
  - defined [24-3](#)
  - enabling [5-23, 24-10](#)
  - notification types [24-10](#)
  - overview [24-1, 24-4](#)
- troubleshooting
  - connectivity problems [34-8, 34-10, 34-11](#)
  - detecting unidirectional links [20-1](#)
  - determining packet forwarding [34-14](#)
  - displaying crash information [34-17](#)
  - PIMv1 and PIMv2 interoperability problems [31-22](#)
  - SFP security and identification [34-8](#)
  - show forward command [34-14](#)
  - troubleshooting (continued)
    - with CiscoWorks [24-4](#)
    - with debug commands [34-13](#)
    - with ping [34-8](#)
    - with system message logging [23-1](#)
    - with traceroute [34-11](#)
- trunk ports
  - and Layer 2 protocol tunneling [13-12](#)
  - configuring [10-21](#)
  - defined [9-3, 10-3](#)
  - encapsulation [10-21, 10-26, 10-27](#)
  - secure MAC addresses on [18-10](#)
- trunks
  - allowed-VLAN list [10-22](#)
  - configuring [10-21, 10-26, 10-27](#)
  - ISL [10-16](#)
  - load sharing
    - setting STP path costs [10-26](#)
    - using STP port priorities [10-25](#)
  - native VLAN for untagged traffic [10-24](#)
  - parallel [10-26](#)
  - pruning-eligible list [10-23](#)
  - to non-DTP device [10-17](#)
  - understanding [10-17](#)
- trusted boundary for QoS [26-46](#)
- trusted port states
  - between QoS domains [26-47](#)
  - classification options [26-6](#)
  - ensuring port security for IP phones [26-46](#)
  - within a QoS domain [26-42](#)
- tunneling
  - 802.1Q [13-2](#)
  - defined [13-1](#)
  - Layer 2 protocol [13-11](#)
- tunnel ports
  - 802.1Q, configuring [13-6](#)
  - defined [10-4](#)
  - described [9-3, 13-2](#)
  - incompatibilities with other features [13-6](#)

twisted-pair Ethernet, detecting unidirectional links [20-1](#)  
 type of service [1-6](#)

---

## U

### UDLD

aggressive mode, described [20-1](#)  
 and autonegotiation [20-1](#)  
 configuration guidelines [20-4](#)  
 default configuration [20-4](#)  
 echoing detection mechanism [20-3](#)  
 enabling  
     globally [20-5](#)  
     per interface [20-5](#)  
 link-detection mechanism [20-1](#)  
 modes of operation [20-1](#)  
 neighbor database [20-2](#)  
 normal mode, described [20-1](#)  
 overview [20-1](#)  
 resetting an interface [20-6](#)  
 status, displaying [20-6](#)  
 unidirectional link, defined [20-2](#)

UDP, configuring [28-14](#)

unauthorized ports with 802.1x [8-4](#)

unequal-cost load balancing, IGRP [28-24](#)

unicast storm control command [18-3](#)

unicast storms [18-1](#)

unicast traffic, blocking [18-6](#)

UniDirectional Link Detection protocol

    See UDLD

UNIX syslog servers

    daemon configuration [23-10](#)

    facilities supported [23-11](#)

    message logging configuration [23-11](#)

unrecognized Type-Length-Value (TLV) support [11-4](#)

upgrading software images

    See downloading

    See release notes [xxxv](#)

UplinkFast

    described [16-4](#)

    enabling [16-13](#)

uploading

    configuration files

        preparing [B-10, B-12, B-16](#)

        reasons for [B-8](#)

        using FTP [B-14](#)

        using RCP [B-17](#)

        using TFTP [B-11](#)

    image files

        preparing [B-21, B-24, B-28](#)

        reasons for [B-19](#)

        using FTP [B-27](#)

        using RCP [B-31](#)

        using TFTP [B-23](#)

User Datagram Protocol

    See UDP

user EXEC mode [2-2](#)

username-based authentication [7-6](#)

---

## V

VCs [30-12](#)

version-dependent transparent mode [11-4](#)

virtual connections

    See VCs

Virtual Private Networks

    See VPNs

virtual router [29-1, 29-2](#)

vlan.dat file [10-5](#)

VLAN 1, disabling on a trunk port [10-22](#)

VLAN 1 minimization [10-22](#)

VLAN ACLs

    See VLAN maps

vlan-assignment response, VMPS [10-28](#)

VLAN configuration

    at bootup [10-8](#)

    saving [10-8](#)

- VLAN configuration mode [2-2, 10-7](#)
- VLAN database
  - and startup configuration file [10-8](#)
  - and VTP [11-1](#)
  - VLAN configuration saved in [10-7](#)
  - VLANs saved in [10-4](#)
- vlan database command [10-7](#)
- VLAN filtering, and SPAN [21-6](#)
- vlan global configuration command [10-7](#)
- VLAN ID
  - customer-side [13-7](#)
  - discovering [5-26](#)
  - mapping [13-7](#)
  - service provider [13-8](#)
- VLAN IDs, number supported [1-4](#)
- VLAN ID translation
  - See VLAN mapping
- VLAN management domain [11-2](#)
- VLAN Management Policy Server
  - See VMPS
- VLAN map entries, order of [25-28](#)
- VLAN mapping
  - 802.1Q traffic [13-9](#)
  - configuring [13-9](#)
  - described [13-7](#)
- VLAN maps
  - applying [25-31](#)
  - common uses for [25-31](#)
  - configuration example [25-32](#)
  - configuration guidelines [25-28](#)
  - configuring [25-27](#)
  - creating [25-28](#)
  - defined [25-2](#)
  - denying access example [25-33](#)
  - denying and permitting packets [25-29](#)
  - displaying [25-38](#)
  - examples [25-33](#)
  - with router ACLs [25-38](#)
- VLAN membership
  - confirming [10-31](#)
  - modes [10-3](#)
- VLAN Query Protocol
  - See VQP
- VLANs
  - adding to VLAN database [10-9](#)
  - aging dynamic addresses [14-9](#)
  - allowed on trunk [10-22](#)
  - and spanning-tree instances [10-3, 10-6, 10-13](#)
  - configuration guidelines
    - extended-range VLANs [10-13](#)
    - normal-range VLANs [10-6](#)
  - configuration options [10-7](#)
  - configuring [10-1](#)
  - configuring IDs 1006 to 4094 [10-13](#)
  - connecting through SVIs [9-5](#)
  - creating [10-9](#)
  - creating in config-vlan mode [10-9](#)
  - creating in VLAN configuration mode [10-10](#)
  - customer numbering in service-provider networks [13-3](#)
  - default configuration [10-8](#)
  - deleting [10-11](#)
  - described [9-2, 10-1](#)
  - displaying [10-16](#)
  - extended-range [10-1, 10-12](#)
  - features [1-4](#)
  - illustrated [10-2](#)
  - internal [10-13](#)
  - limiting source traffic
    - with RSPAN [21-22](#)
    - with SPAN [21-15](#)
  - modifying [10-9](#)
  - native, configuring [10-24](#)
  - normal-range [10-1, 10-4](#)
  - number supported [1-4](#)
  - parameters [10-5](#)
  - port membership modes [10-3](#)

## VLANs (continued)

- static-access ports [10-11](#)
- STP and 802.1Q trunks [14-10](#)
- supported [10-3](#)
- Token Ring [10-6](#)
- traffic between [10-2](#)
- VLAN-bridge STP [14-11, 33-1](#)
- VTP modes [11-3](#)

## VLAN Trunking Protocol

See VTP

VLAN trunks [10-16, 10-17](#)

## VMPS

- administering [10-33](#)
- configuration example [10-34](#)
- configuration guidelines [10-29](#)
- default configuration [10-29](#)
- description [10-28](#)
- dynamic port membership
  - described [10-29](#)
  - reconfirming [10-32](#)
  - troubleshooting [10-33](#)
- entering server address [10-30](#)
- mapping MAC addresses to VLANs [10-28](#)
- monitoring [10-33](#)
- reconfirmation interval, changing [10-32](#)
- reconfirming membership [10-31](#)
- retry count, changing [10-32](#)

voice-over-IP [12-1](#)

## voice VLAN

- Cisco 7960 phone, port connections [12-1](#)
- configuration guidelines [12-3](#)
- configuring IP phones for data traffic
  - override CoS of incoming frame [12-5](#)
  - trust CoS priority of incoming frame [12-5](#)
- configuring ports for voice traffic in
  - 802.1P priority tagged frames [12-5](#)
  - 802.1Q frames [12-4](#)

## voice VLAN (continued)

- connecting to an IP phone [12-4](#)
- default configuration [12-3](#)
- described [12-1](#)
- displaying [12-6](#)

## VPN routing and forwarding table

See VRF

## VPNs

- and multi-VRF CE [28-80](#)
- benefits [30-3](#)
- configuring [30-8](#)
- configuring routing sessions [28-80](#)
- described [30-2](#)
- forwarding in [28-78](#)
- in service provider networks [28-76](#)
- IPv4 prefix [30-5](#)
- MPLS [30-6](#)
- number supported [30-6](#)
- routes [28-76, 30-2, 30-5](#)

## VPN services

- Layer 2 [1-4](#)
- Layer 3 [1-5](#)

VQP [10-28](#)

## VRF

- configuration [30-8](#)
- defining [28-78](#)
- elements [30-3](#)
- in MPLS VPNs [30-2](#)
- tables [28-76](#)

## VTP

- adding a client to a domain [11-15](#)
- advertisements [10-20, 11-3](#)
- and extended-range VLANs [11-1](#)
- and normal-range VLANs [11-2](#)
- client mode, configuring [11-11](#)



## VTP (continued)

- configuration
  - global configuration mode [11-7](#)
  - guidelines [11-8](#)
  - privileged EXEC mode [11-7](#)
  - requirements [11-9](#)
  - saving [11-7](#)
  - VLAN configuration mode [11-7](#)
- configuration mode options [11-7](#)
- configuration requirements [11-9](#)
- configuration revision number
  - guideline [11-15](#)
  - resetting [11-15](#)
- configuring
  - client mode [11-11](#)
  - server mode [11-9](#)
  - transparent mode [11-12](#)
- consistency checks [11-4](#)
- default configuration [11-6](#)
- described [11-1](#)
- disabling [11-12](#)
- domain names [11-8](#)
- domains [11-2](#)
- Layer 2 protocol tunneling [13-11](#)
- modes
  - client [11-3, 11-11](#)
  - server [11-3, 11-9](#)
  - transitions [11-3](#)
  - transparent [11-3, 11-12](#)
- monitoring [11-16](#)
- passwords [11-8](#)
- pruning
  - disabling [11-14](#)
  - enabling [11-14](#)
  - examples [11-5](#)
  - overview [11-4](#)
- pruning-eligible list, changing [10-23](#)
- server mode, configuring [11-9](#)
- statistics [11-16](#)

## VTP (continued)

- Token Ring support [11-4](#)
- transparent mode, configuring [11-12](#)
- using [11-1](#)
- version, guidelines [11-8](#)
- version 1 [11-4](#)
- version 2
  - configuration guidelines [11-8](#)
  - disabling [11-13](#)
  - enabling [11-13](#)
  - overview [11-4](#)

---

W

## Weighted Random Early Detection

- See WRED

## weighted tail drop

- See WTD

## WRED

- configuring DSCP-based [26-89](#)
- configuring IP precedence-based [26-93](#)
- described [26-28](#)

## WTD

- described [26-13](#)
- setting thresholds
  - egress queue-sets [26-69](#)
  - ingress queues [26-65](#)

---

X

- XMODEM protocol [34-2](#)

