



INDEX

A

- abbreviating commands [2-3](#)
- AC (command switch) [6-9](#)
- access-class command [31-17](#)
- access control entries
 - See ACEs
- access-denied response, VMPS [13-23](#)
- access groups
 - Layer 3 [31-18](#)
- access groups, applying IPv4 ACLs to interfaces [31-18](#)
- accessing
 - clusters, switch [6-12](#)
 - command switches [6-10](#)
 - member switches [6-12](#)
 - switch clusters [6-12](#)
- accessing stack members [7-21](#)
- access lists
 - See ACLs
- access ports
 - in switch clusters [6-8](#)
- access ports, defined [12-3](#)
- accounting
 - with 802.1x [10-50](#)
 - with IEEE 802.1x [10-16](#)
 - with RADIUS [9-35](#)
 - with TACACS+ [9-11, 9-17](#)
- ACEs
 - and QoS [33-8](#)
 - defined [31-2](#)
 - Ethernet [31-2](#)
 - IP [31-2](#)

ACLs

- ACEs [31-2](#)
- any keyword [31-11](#)
- applying
 - time ranges to [31-15](#)
 - to an interface [31-17](#)
 - to QoS [33-8](#)
- classifying traffic for QoS [33-49](#)
- comments in [31-16](#)
- compiling [31-20](#)
- defined [31-1, 31-7](#)
- examples of [31-20, 33-49](#)
- extended IP, configuring for QoS classification [33-50](#)
- extended IPv4
 - creating [31-9](#)
 - matching criteria [31-7](#)
- hardware and software handling [31-19](#)
- host keyword [31-11](#)
- IP
 - creating [31-7](#)
 - fragments and QoS guidelines [33-40](#)
 - implicit deny [31-9, 31-12, 31-14](#)
 - implicit masks [31-9](#)
 - matching criteria [31-7](#)
 - undefined [31-19](#)
- IPv4
 - applying to interfaces [31-17](#)
 - creating [31-7](#)
 - matching criteria [31-7](#)
 - named [31-13](#)
 - numbers [31-7](#)
 - terminal lines, setting on [31-17](#)

- unsupported features 31-6
 - MAC extended 31-22, 33-51
 - matching 31-7, 31-19
 - monitoring 31-24
 - named, IPv4 31-13
 - number per QoS class map 33-40
 - port 31-2
 - precedence of 31-2
 - QoS 33-8, 33-49
 - resequencing entries 31-13
 - router 31-2
 - standard IP, configuring for QoS classification 33-49
 - standard IPv4
 - creating 31-8
 - matching criteria 31-7
 - support for 1-11
 - support in hardware 31-19
 - time ranges 31-15
 - types supported 31-2
 - unsupported features, IPv4 31-6
- active link 19-4, 19-5, 19-6
- active links 19-2
- active traffic monitoring, IP SLAs 32-1
- address aliasing 21-2
- addresses
 - displaying the MAC address table 5-24
 - dynamic
 - accelerated aging 16-9
 - changing the aging time 5-15
 - default aging 16-9
 - defined 5-13
 - learning 5-14
 - removing 5-16
 - IPv6 35-2
 - MAC, discovering 5-24
 - multicast, STP address management 16-9
 - static
 - adding and removing 5-20
 - defined 5-13
 - address resolution 5-24
 - Address Resolution Protocol
 - See ARP
 - advertisements
 - CDP 25-1
 - LLDP 26-2
 - VTP 13-15, 14-3, 14-4
 - aggregatable global unicast addresses 35-3
 - aggregated ports
 - See EtherChannel
 - aggregate policers 33-59
 - aggregate policing 1-14
 - aging, accelerating 16-9
 - aging time
 - accelerated
 - for MSTP 17-24
 - for STP 16-9, 16-23
 - MAC address table 5-15
 - maximum
 - for MSTP 17-24, 17-25
 - for STP 16-23, 16-24
 - alarms, RMON 28-3
 - allowed-VLAN list 13-17
 - ARP
 - defined 1-6, 5-24
 - table
 - address resolution 5-24
 - managing 5-24
 - attributes, RADIUS
 - vendor-proprietary 9-38
 - vendor-specific 9-36
 - attribute-value pairs 10-13, 10-16, 10-21, 10-22
 - authentication
 - local mode with AAA 9-40
 - open1x 10-30
 - RADIUS
 - key 9-28
 - login 9-30

- TACACS+
 - defined [9-11](#)
 - key [9-13](#)
 - login [9-14](#)
 See also port-based authentication
 - authentication compatibility with Catalyst 6000 switches [10-8](#)
 - authentication failed VLAN
 - See restricted VLAN
 - authentication manager
 - CLI commands [10-9](#)
 - compatibility with older 802.1x CLI commands [10-9 to ??](#)
 - overview [10-7](#)
 - authoritative time source, described [5-3](#)
 - authorization
 - with RADIUS [9-34](#)
 - with TACACS+ [9-11, 9-16](#)
 - authorized ports with IEEE 802.1x [10-10](#)
 - autoconfiguration [3-3](#)
 - auto enablement [10-31](#)
 - automatic advise (auto-advise) in switch stacks [7-11](#)
 - automatic copy (auto-copy) in switch stacks [7-10](#)
 - automatic discovery
 - considerations
 - beyond a noncandidate device [6-8](#)
 - brand new switches [6-8](#)
 - connectivity [6-5](#)
 - different VLANs [6-7](#)
 - management VLANs [6-7](#)
 - non-CDP-capable devices [6-6](#)
 - noncluster-capable devices [6-6](#)
 - in switch clusters [6-5](#)
 - See also CDP
 - automatic extraction (auto-extract) in switch stacks [7-10](#)
 - automatic QoS
 - See QoS
 - automatic recovery, clusters [6-9](#)
 - See also HSRP
 - automatic upgrades (auto-upgrade) in switch stacks [7-10](#)
 - auto-MDIX
 - configuring [12-29](#)
 - described [12-29](#)
 - autonegotiation
 - duplex mode [1-4](#)
 - interface configuration guidelines [12-26](#)
 - mismatches [38-12](#)
 - Auto-QoS video devices [1-15](#)
 - autosensing, port speed [1-4](#)
 - auxiliary VLAN
 - See voice VLAN
 - availability, features [1-8](#)
-
- ## B
- BackboneFast
 - described [18-7](#)
 - disabling [18-17](#)
 - enabling [18-17](#)
 - support for [1-8](#)
 - backup interfaces
 - See Flex Links
 - backup links [19-2](#)
 - banners
 - configuring
 - login [5-13](#)
 - message-of-the-day login [5-12](#)
 - default configuration [5-11](#)
 - when displayed [5-11](#)
 - Berkeley r-tools replacement [9-52](#)
 - binding database
 - DHCP snooping
 - See DHCP snooping binding database
 - bindings
 - DHCP snooping database [20-6](#)
 - IP source guard [20-13](#)
 - binding table, DHCP snooping
 - See DHCP snooping binding database

- blocking packets 23-7
 - booting
 - boot loader, function of 3-2
 - boot process 3-1
 - manually 3-19
 - specific image 3-19
 - boot loader
 - accessing 3-20
 - described 3-2
 - environment variables 3-20
 - prompt 3-20
 - trap-door mechanism 3-2
 - BPDUs
 - error-disabled state 18-2
 - filtering 18-3
 - RSTP format 17-12
 - BPDUs filtering
 - described 18-3
 - disabling 18-15
 - enabling 18-15
 - support for 1-9
 - BPDUs guard
 - described 18-2
 - disabling 18-14
 - enabling 18-14
 - support for 1-9
 - bridge protocol data unit
 - See BPDUs
 - broadcast storm-control command 23-4
 - broadcast storms 23-1
-
- C**
- cables, monitoring for unidirectional links 24-1
 - candidate switch
 - automatic discovery 6-5
 - defined 6-4
 - requirements 6-4
 - See also command switch, cluster standby group, and member switch
 - Catalyst 6000 switches
 - authentication compatibility 10-8
 - CA trustpoint
 - configuring 9-49
 - defined 9-46
 - CDP
 - and trusted boundary 33-45
 - automatic discovery in switch clusters 6-5
 - configuring 25-2
 - default configuration 25-2
 - defined with LLDP 26-1
 - described 25-1
 - disabling for routing device 25-4
 - enabling and disabling
 - on an interface 25-4
 - on a switch 25-4
 - monitoring 25-5
 - overview 25-1
 - power negotiation extensions 12-5
 - support for 1-6
 - switch stack considerations 25-2
 - transmission timer and holdtime, setting 25-3
 - updates 25-3
 - CGMP
 - as IGMP snooping learning method 21-9
 - joining multicast group 21-3
 - CipherSuites 9-48
 - Cisco 7960 IP Phone 15-1
 - Cisco Discovery Protocol
 - See CDP
 - Cisco intelligent power management 12-5
 - Cisco IOS File System
 - See IFS
 - Cisco IOS IP SLAs 32-1
 - Cisco Secure ACS
 - attribute-value pairs for downloadable ACLs 10-22
 - attribute-value pairs for redirect URL 10-21

- Cisco Secure ACS configuration guide [10-60](#)
- CiscoWorks 2000 [1-6, 30-5](#)
- CISP [10-31](#)
- CIST regional root
 - See MSTP
- CIST root
 - See MSTP
- civic location [26-3](#)
- class maps for QoS
 - configuring [33-52](#)
 - described [33-8](#)
 - displaying [33-80](#)
- class of service
 - See CoS
- clearing interfaces [12-38](#)
- CLI
 - abbreviating commands [2-3](#)
 - command modes [2-1](#)
 - configuration logging [2-4](#)
 - described [1-6](#)
 - editing features
 - enabling and disabling [2-6](#)
 - keystroke editing [2-7](#)
 - wrapped lines [2-8](#)
 - error messages [2-4](#)
 - filtering command output [2-9](#)
 - getting help [2-3](#)
 - history
 - changing the buffer size [2-5](#)
 - described [2-5](#)
 - disabling [2-6](#)
 - recalling commands [2-5](#)
 - managing clusters [6-15](#)
 - no and default forms of commands [2-4](#)
- Client Information Signalling Protocol
 - See CISP
- client mode, VTP [14-3](#)
- clock
 - See system clock
- clusters, switch
 - accessing [6-12](#)
 - automatic discovery [6-5](#)
 - automatic recovery [6-9](#)
 - benefits [1-2](#)
 - compatibility [6-4](#)
 - described [6-1](#)
 - LRE profile considerations [6-15](#)
 - managing
 - through CLI [6-15](#)
 - through SNMP [6-16](#)
 - planning [6-4](#)
 - planning considerations
 - automatic discovery [6-5](#)
 - automatic recovery [6-9](#)
 - CLI [6-15](#)
 - host names [6-12](#)
 - IP addresses [6-12](#)
 - LRE profiles [6-15](#)
 - passwords [6-13](#)
 - RADIUS [6-15](#)
 - SNMP [6-13, 6-16](#)
 - switch stacks [6-13](#)
 - TACACS+ [6-15](#)

See also candidate switch, command switch, cluster standby group, member switch, and standby command switch
- cluster standby group
 - automatic recovery [6-11](#)
 - considerations [6-10](#)
 - defined [6-2](#)
 - requirements [6-3](#)
 - virtual IP address [6-10](#)

See also HSRP

CNS 1-6

Configuration Engine

configID, deviceID, hostname 4-3

configuration service 4-2

described 4-1

event service 4-3

embedded agents

described 4-5

enabling automated configuration 4-6

enabling configuration agent 4-9

enabling event agent 4-7

management functions 1-6

CoA Request Commands 9-23

Coarse Wave Division Multiplexer

See CWDM SFPs

command-line interface

See CLI

command modes 2-1

commands

abbreviating 2-3

no and default 2-4

commands, setting privilege levels 9-8

command switch

accessing 6-10

active (AC) 6-9

configuration conflicts 38-12

defined 6-2

passive (PC) 6-9

password privilege levels 6-16

priority 6-9

recovery

from command-switch failure 6-9, 38-8

from lost member connectivity 38-12

redundant 6-9

replacing

with another switch 38-10

with cluster member 38-9

requirements 6-3

standby (SC) 6-9

See also candidate switch, cluster standby group, member switch, and standby command switch

community strings

configuring 6-13, 30-8

for cluster switches 30-4

in clusters 6-13

overview 30-4

SNMP 6-13

compatibility, feature 23-12

compatibility, software

See stacks, switch

config.text 3-18

configurable leave timer, IGMP 21-6

configuration, initial

defaults 1-16

Express Setup 1-2

configuration changes, logging 29-11

configuration conflicts, recovering from lost member connectivity 38-12

configuration examples, network 1-19

configuration files

archiving A-20

clearing the startup configuration A-19

creating using a text editor A-10

default name 3-18

deleting a stored configuration A-19

described A-8

downloading

automatically 3-18

preparing A-10, A-13, A-16

reasons for A-8

using FTP A-13

using RCP A-17

using TFTP A-11

guidelines for creating and using A-9

guidelines for replacing and rolling back A-21

invalid combinations when copying A-5

limiting TFTP server access 30-17

obtaining with DHCP 3-9

- password recovery disable considerations 9-5
 - replacing a running configuration A-19, A-20
 - rolling back a running configuration A-19, A-21
 - specifying the filename 3-18
 - system contact and location information 30-17
 - types and location A-10
 - uploading
 - preparing A-10, A-13, A-16
 - reasons for A-8
 - using FTP A-15
 - using RCP A-18
 - using TFTP A-12
 - configuration logger 29-11
 - configuration logging 2-4
 - configuration replacement A-19
 - configuration rollback A-19, A-20
 - configuration settings, saving 3-15
 - configure terminal command 12-16
 - configuring 802.1x user distribution 10-56
 - configuring port-based authentication violation modes 10-40 to 10-41
 - configuring small-frame arrival rate 23-5
 - conflicts, configuration 38-12
 - connections, secure remote 9-42
 - connectivity problems 38-14, 38-15, 38-17
 - consistency checks in VTP Version 2 14-5
 - console port, connecting to 2-10
 - control protocol, IP SLAs 32-4
 - corrupted software, recovery steps with Xmodem 38-2
 - CoS
 - in Layer 2 frames 33-2
 - override priority 15-6
 - trust priority 15-6
 - CoS input queue threshold map for QoS 33-16
 - CoS output queue threshold map for QoS 33-19
 - CoS-to-DSCP map for QoS 33-62
 - counters, clearing interface 12-38
 - CPU utilization, troubleshooting 38-27
 - crashinfo file 38-23
 - critical authentication, IEEE 802.1x 10-54
 - critical VLAN 10-24
 - cross-stack EtherChannel
 - configuration guidelines 37-13
 - described 37-3
 - illustration 37-4
 - support for 1-8
 - cross-stack UplinkFast, STP
 - described 18-5
 - disabling 18-17
 - enabling 18-16
 - fast-convergence events 18-7
 - Fast Uplink Transition Protocol 18-6
 - normal-convergence events 18-7
 - support for 1-8
 - cryptographic software image
 - SSH 9-41
 - SSL 9-46
 - switch stack considerations 7-14
 - customizable web pages, web-based authentication 11-6
 - CWDM SFPs 1-23
-
- ## D
- DAACL
 - See downloadable ACL
 - daylight saving time 5-7
 - debugging
 - enabling all system diagnostics 38-20
 - enabling for a specific feature 38-20
 - redirecting error message output 38-21
 - using commands 38-19
 - default commands 2-4
 - default configuration
 - 802.1x 10-35
 - auto-QoS 33-22
 - banners 5-11
 - CDP 25-2
 - DHCP 20-8

- DHCP option 82 [20-8](#)
- DHCP snooping [20-8](#)
- DHCP snooping binding database [20-8](#)
- DNS [5-10](#)
- dynamic ARP inspection [22-5](#)
- EtherChannel [37-11](#)
- Ethernet interfaces [12-23](#)
- Flex Links [19-7, 19-8](#)
- IGMP filtering [21-24](#)
- IGMP snooping [21-7, 36-6](#)
- IGMP throttling [21-24](#)
- initial switch information [3-3](#)
- IP SLAs [32-5](#)
- IP source guard [20-15](#)
- IPv6 [35-6](#)
- Layer 2 interfaces [12-23](#)
- LLDP [26-5](#)
- MAC address table [5-15](#)
- MAC address-table move update [19-8](#)
- MSTP [17-14](#)
- MVR [21-19](#)
- optional spanning-tree configuration [18-12](#)
- password and privilege level [9-2](#)
- RADIUS [9-27](#)
- RMON [28-3](#)
- RSPAN [27-10](#)
- SDM template [8-3](#)
- SNMP [30-7](#)
- SPAN [27-10](#)
- SSL [9-48](#)
- standard QoS [33-38](#)
- STP [16-13](#)
- switch stacks [7-17](#)
- system message logging [29-4](#)
- system name and prompt [5-9](#)
- TACACS+ [9-13](#)
- UDLD [24-4](#)
- VLAN, Layer 2 Ethernet interfaces [13-15](#)
- VLANs [13-7](#)
- VMPS [13-24](#)
- voice VLAN [15-3](#)
- VTP [14-9](#)
- default gateway [3-15](#)
- default web-based authentication configuration
 - 802.1X [11-9](#)
- deleting VLANs [13-9](#)
- denial-of-service attack [23-1](#)
- description command [12-34](#)
- designing your network, examples [1-19](#)
- destination addresses
 - in IPv4 ACLs [31-10](#)
- destination-IP address-based forwarding, EtherChannel [37-9](#)
- destination-MAC address forwarding, EtherChannel [37-9](#)
- detecting indirect link failures, STP [18-8](#)
- device [A-24](#)
- device discovery protocol [25-1, 26-1](#)
- device manager
 - benefits [1-2](#)
 - described [1-2, 1-5](#)
 - in-band management [1-7](#)
 - upgrading a switch [A-24](#)
- DHCP
 - enabling
 - relay agent [20-10](#)
- DHCP-based autoconfiguration
 - client request message exchange [3-4](#)
 - configuring
 - client side [3-3](#)
 - DNS [3-8](#)
 - relay device [3-8](#)
 - server side [3-6](#)
 - TFTP server [3-7](#)
 - example [3-9](#)
 - lease options
 - for IP address information [3-6](#)
 - for receiving the configuration file [3-6](#)

- overview [3-3](#)
- relationship to BOOTP [3-3](#)
- relay support [1-6](#)
- support for [1-6](#)
- DHCP-based autoconfiguration and image update
 - configuring [3-11 to 3-14](#)
 - understanding [3-5 to 3-6](#)
- DHCP binding database
 - See DHCP snooping binding database
- DHCP binding table
 - See DHCP snooping binding database
- DHCP option 82
 - circuit ID suboption [20-5](#)
 - configuration guidelines [20-8](#)
 - default configuration [20-8](#)
 - displaying [20-13](#)
 - overview [20-3](#)
 - packet format, suboption
 - circuit ID [20-5](#)
 - remote ID [20-5](#)
 - remote ID suboption [20-5](#)
- DHCP server port-based address allocation
 - configuration guidelines [20-22](#)
 - default configuration [20-21](#)
 - described [20-21](#)
 - displaying [20-24](#)
 - enabling [20-22](#)
 - reserved addresses [20-22](#)
- DHCP server port-based address assignment
 - support for [1-6](#)
- DHCP snooping
 - accepting untrusted packets form edge switch [20-3, 20-10](#)
 - binding database
 - See DHCP snooping binding database
 - configuration guidelines [20-8](#)
 - default configuration [20-8](#)
 - displaying binding tables [20-13](#)
 - message exchange process [20-4](#)
 - option 82 data insertion [20-3](#)
 - trusted interface [20-2](#)
 - untrusted interface [20-2](#)
 - untrusted messages [20-2](#)
 - DHCP snooping binding database
 - adding bindings [20-12](#)
 - binding entries, displaying [20-13](#)
 - binding file
 - format [20-7](#)
 - location [20-6](#)
 - bindings [20-6](#)
 - clearing agent statistics [20-12](#)
 - configuration guidelines [20-9](#)
 - configuring [20-12](#)
 - default configuration [20-8](#)
 - deleting
 - binding file [20-12](#)
 - bindings [20-13](#)
 - database agent [20-12](#)
 - described [20-6](#)
 - displaying [20-13](#)
 - displaying status and statistics [20-13](#)
 - enabling [20-12](#)
 - entry [20-6](#)
 - renewing database [20-12](#)
 - resetting
 - delay value [20-12](#)
 - timeout value [20-12](#)
 - DHCP snooping binding table
 - See DHCP snooping binding database
- Differentiated Services architecture, QoS [33-2](#)
- Differentiated Services Code Point [33-2](#)
- directed unicast requests [1-6](#)
- directories
 - changing [A-4](#)
 - creating and removing [A-4](#)
 - displaying the working [A-4](#)
- discovery, clusters
 - See automatic discovery

- DNS
 - and DHCP-based autoconfiguration [3-8](#)
 - default configuration [5-10](#)
 - displaying the configuration [5-11](#)
 - in IPv6 [35-3](#)
 - overview [5-9](#)
 - setting up [5-10](#)
 - support for [1-6](#)
- domain names
 - DNS [5-9](#)
 - VTP [14-9](#)
- Domain Name System
 - See DNS
- downloadable ACL [10-20, 10-22, 10-60](#)
- downloading
 - configuration files
 - preparing [A-10, A-13, A-16](#)
 - reasons for [A-8](#)
 - using FTP [A-13](#)
 - using RCP [A-17](#)
 - using TFTP [A-11](#)
 - image files
 - deleting old image [A-28](#)
 - preparing [A-26, A-30, A-34](#)
 - reasons for [A-24](#)
 - using CMS [1-2](#)
 - using FTP [A-31](#)
 - using HTTP [1-2, A-24](#)
 - using RCP [A-35](#)
 - using TFTP [A-27](#)
 - using the device manager or Network Assistant [A-24](#)
- DRP
 - support for [1-15](#)
- DSCP [1-13, 33-2](#)
- DSCP input queue threshold map for QoS [33-16](#)
- DSCP output queue threshold map for QoS [33-19](#)
- DSCP-to-CoS map for QoS [33-65](#)
- DSCP-to-DSCP-mutation map for QoS [33-66](#)
- DSCP transparency [33-46](#)
- DTP [1-9, 13-13](#)
- dual-action detection [37-6](#)
- dual IPv4 and IPv6 templates [35-4, 35-5](#)
- dual protocol stacks
 - IPv4 and IPv6 [35-5](#)
 - SDM templates supporting [35-5](#)
- dual-purpose uplinks
 - defined [12-4](#)
 - LEDs [12-5](#)
 - link selection [12-4, 12-24](#)
 - setting the type [12-24](#)
- dynamic access ports
 - characteristics [13-4](#)
 - configuring [13-26](#)
 - defined [12-3](#)
- dynamic addresses
 - See addresses
- dynamic ARP inspection
 - ARP cache poisoning [22-1](#)
 - ARP requests, described [22-1](#)
 - ARP spoofing attack [22-1](#)
 - clearing
 - log buffer [22-16](#)
 - statistics [22-16](#)
 - configuration guidelines [22-6](#)
 - configuring
 - ACLs for non-DHCP environments [22-9](#)
 - in DHCP environments [22-7](#)
 - log buffer [22-13](#)
 - rate limit for incoming ARP packets [22-4, 22-11](#)
 - default configuration [22-5](#)
 - denial-of-service attacks, preventing [22-11](#)
 - described [22-1](#)
 - DHCP snooping binding database [22-2](#)
 - displaying
 - ARP ACLs [22-15](#)
 - configuration and operating state [22-15](#)

- log buffer [22-16](#)
- statistics [22-16](#)
- trust state and rate limit [22-15](#)
- error-disabled state for exceeding rate limit [22-4](#)
- function of [22-2](#)
- interface trust states [22-3](#)
- log buffer
 - clearing [22-16](#)
 - configuring [22-13](#)
 - displaying [22-16](#)
- logging of dropped packets, described [22-5](#)
- man-in-the middle attack, described [22-2](#)
- network security issues and interface trust states [22-3](#)
- priority of ARP ACLs and DHCP snooping entries [22-4](#)
- rate limiting of ARP packets
 - configuring [22-11](#)
 - described [22-4](#)
 - error-disabled state [22-4](#)
- statistics
 - clearing [22-16](#)
 - displaying [22-16](#)
- validation checks, performing [22-12](#)
- dynamic auto trunking mode [13-14](#)
- dynamic desirable trunking mode [13-14](#)
- Dynamic Host Configuration Protocol
 - See DHCP-based autoconfiguration
- dynamic port VLAN membership
 - described [13-24](#)
 - reconfirming [13-26](#)
 - troubleshooting [13-28](#)
 - types of connections [13-26](#)
- Dynamic Trunking Protocol
 - See DTP

E

- editing features

- enabling and disabling [2-6](#)
- keystrokes used [2-7](#)
- wrapped lines [2-8](#)
- elections
 - See stack master
- ELIN location [26-3](#)
- enable password [9-3](#)
- enable secret password [9-3](#)
- encryption, CipherSuite [9-48](#)
- encryption for passwords [9-3](#)
- environment variables, function of [3-21](#)
- error-disabled state, BPDU [18-2](#)
- error messages during command entry [2-4](#)
- EtherChannel
 - automatic creation of [37-5, 37-7](#)
 - channel groups
 - binding physical and logical interfaces [37-4](#)
 - numbering of [37-4](#)
 - configuration guidelines [37-11](#)
 - configuring Layer 2 interfaces [37-13](#)
 - default configuration [37-11](#)
 - described [37-2](#)
 - displaying status [37-20](#)
 - forwarding methods [37-8, 37-15](#)
 - IEEE 802.3ad, described [37-7](#)
 - interaction
 - with STP [37-12](#)
 - with VLANs [37-12](#)
- LACP
 - described [37-7](#)
 - displaying status [37-20](#)
 - hot-standby ports [37-18](#)
 - interaction with other features [37-8](#)
 - modes [37-7](#)
 - port priority [37-19](#)
 - system priority [37-18](#)
- load balancing [37-8, 37-15](#)
- PAgP
 - aggregate-port learners [37-16](#)

- compatibility with Catalyst 1900 37-17
 - described 37-5
 - displaying status 37-20
 - interaction with other features 37-7
 - interaction with virtual switches 37-6
 - learn method and priority configuration 37-16
 - modes 37-6
 - support for 1-4
 - with dual-action detection 37-6
 - port-channel interfaces
 - described 37-4
 - numbering of 37-4
 - port groups 12-4
 - stack changes, effects of 37-10
 - support for 1-4
 - EtherChannel guard
 - described 18-10
 - disabling 18-18
 - enabling 18-17
 - Ethernet management port
 - active link 12-21
 - and routing 12-21
 - and TFTP 12-22
 - configuring 12-22
 - default setting 12-21
 - described 12-20
 - for network management 12-20
 - specifying 12-22
 - supported features 12-21
 - unsupported features 12-22
 - Ethernet management port, internal
 - and routing 12-21
 - unsupported features 12-22
 - Ethernet VLANs
 - adding 13-8
 - defaults and ranges 13-7
 - modifying 13-8
 - EUI 35-3
 - events, RMON 28-3
 - examples
 - network configuration 1-19
 - expedite queue for QoS 33-79
 - Express Setup 1-2
 - See also getting started guide
 - extended crashinfo file 38-23
 - extended-range VLANs
 - configuration guidelines 13-11
 - configuring 13-10
 - creating 13-11
 - defined 13-1
 - extended system ID
 - MSTP 17-18
 - STP 16-4, 16-16
 - extended universal identifier
 - See EUI
 - Extensible Authentication Protocol over LAN 10-1
-
- F**
- fa0 interface 1-7
 - Fa0 port
 - See Ethernet management port
 - failover support 1-8
 - Fast Convergence 19-3
 - fastethernet0 port
 - See Ethernet management port
 - Fast Uplink Transition Protocol 18-6
 - features, incompatible 23-12
 - fiber-optic, detecting unidirectional links 24-1
 - files
 - basic crashinfo
 - description 38-23
 - location 38-23
 - copying A-5
 - crashinfo, description 38-23
 - deleting A-5
 - displaying the contents of A-8
 - extended crashinfo

- description [38-23](#)
 - location [38-23](#)
 - tar
 - creating [A-6](#)
 - displaying the contents of [A-6](#)
 - extracting [A-7](#)
 - image file format [A-25](#)
 - file system
 - displaying available file systems [A-2](#)
 - displaying file information [A-3](#)
 - local file system names [A-1](#)
 - network file system names [A-5](#)
 - setting the default [A-3](#)
 - filtering
 - non-IP traffic [31-22](#)
 - show and more command output [2-9](#)
 - filtering show and more command output [2-9](#)
 - filters, IP
 - See ACLs, IP
 - flash device, number of [A-1](#)
 - flexible authentication ordering
 - configuring [10-63](#)
 - overview [10-30](#)
 - Flex Link Multicast Fast Convergence [19-3](#)
 - Flex Links
 - configuration guidelines [19-8](#)
 - configuring [19-8, 19-9](#)
 - configuring preferred VLAN [19-12](#)
 - configuring VLAN load balancing [19-10](#)
 - default configuration [19-7](#)
 - description [19-2](#)
 - link load balancing [19-2](#)
 - monitoring [19-14](#)
 - VLANs [19-2](#)
 - flooded traffic, blocking [23-8](#)
 - flow-based packet classification [1-13](#)
 - flowcharts
 - QoS classification [33-7](#)
 - QoS egress queueing and scheduling [33-18](#)
 - QoS ingress queueing and scheduling [33-15](#)
 - QoS policing and marking [33-11](#)
 - flowcontrol
 - configuring [12-28](#)
 - described [12-28](#)
 - forward-delay time
 - MSTP [17-24](#)
 - STP [16-23](#)
 - FTP
 - configuration files
 - downloading [A-13](#)
 - overview [A-12](#)
 - preparing the server [A-13](#)
 - uploading [A-15](#)
 - image files
 - deleting old image [A-32](#)
 - downloading [A-31](#)
 - preparing the server [A-30](#)
 - uploading [A-32](#)
-
- ## G
- general query [19-5](#)
 - Generating IGMP Reports [19-4](#)
 - get-bulk-request operation [30-4](#)
 - get-next-request operation [30-3, 30-5](#)
 - get-request operation [30-3, 30-4, 30-5](#)
 - get-response operation [30-4](#)
 - Gigabit modules
 - See SFPs
 - global configuration mode [2-2](#)
 - global leave, IGMP [21-13](#)
 - guest VLAN and 802.1x [10-22](#)
 - guide mode [1-2](#)
 - GUIs
 - See device manager and Network Assistant

H

hello time

 MSTP [17-23](#)

 STP [16-22](#)

help, for the command line [2-3](#)

HFTM space [38-26](#)

history

 changing the buffer size [2-5](#)

 described [2-5](#)

 disabling [2-6](#)

 recalling commands [2-5](#)

history table, level and number of syslog messages [29-10](#)

host names, in clusters [6-12](#)

hosts, limit on dynamic ports [13-28](#)

HP OpenView [1-6](#)

HQATM space [38-26](#)

HSRP

 automatic cluster recovery [6-11](#)

 cluster standby group considerations [6-10](#)

 See also clusters, cluster standby group, and standby command switch

HTTP over SSL

 see HTTPS

HTTPS [9-46](#)

 configuring [9-50](#)

 self-signed certificate [9-47](#)

HTTP secure server [9-46](#)

Hulc Forwarding TCAM Manager

 See HFTM space

Hulc QoS/ACL TCAM Manager

 See HQATM space

I

ICMP

 IPv6 [35-3](#)

 time-exceeded messages [38-17](#)

 traceroute and [38-17](#)

ICMP ping

 executing [38-14](#)

 overview [38-14](#)

ICMPv6 [35-3](#)

IDS appliances

 and ingress RSPAN [27-21](#)

 and ingress SPAN [27-14](#)

IEEE 802.1D

 See STP

IEEE 802.1p [15-1](#)

IEEE 802.1Q

 and trunk ports [12-3](#)

 configuration limitations [13-14](#)

 encapsulation [13-13](#)

 native VLAN for untagged traffic [13-18](#)

IEEE 802.1s

 See MSTP

IEEE 802.1w

 See RSTP

IEEE 802.1x

 See port-based authentication

IEEE 802.3ad

 See EtherChannel

IEEE 802.3ad, PoE+ [1-15, 12-6](#)

IEEE 802.3af

 See PoE

IEEE 802.3x flow control [12-28](#)

ifIndex values, SNMP [30-6](#)

IFS [1-7](#)

IGMP

 configurable leave timer

 described [21-6](#)

 enabling [21-11](#)

 flooded multicast traffic

 controlling the length of time [21-12](#)

 disabling on an interface [21-13](#)

 global leave [21-13](#)

 query solicitation [21-13](#)

 recovering from flood mode [21-13](#)

- joining multicast group 21-3
- join messages 21-3
- leave processing, enabling 21-10, 36-9
- leaving multicast group 21-5
- queries 21-4
- report suppression
 - described 21-6
 - disabling 21-15, 36-11
- supported versions 21-3
- support for 1-4
- IGMP filtering
 - configuring 21-24
 - default configuration 21-24
 - described 21-23
 - monitoring 21-28
 - support for 1-5
- IGMP groups
 - configuring filtering 21-27
 - setting the maximum number 21-26
- IGMP Immediate Leave
 - configuration guidelines 21-11
 - described 21-5
 - enabling 21-10
- IGMP profile
 - applying 21-25
 - configuration mode 21-24
 - configuring 21-25
- IGMP snooping
 - and address aliasing 21-2
 - and stack changes 21-6
 - configuring 21-7
 - default configuration 21-7, 36-6
 - definition 21-2
 - enabling and disabling 21-7, 36-7
 - global configuration 21-7
 - Immediate Leave 21-5
 - in the switch stack 21-6
 - method 21-8
 - monitoring 21-15, 36-11
 - querier
 - configuration guidelines 21-14
 - configuring 21-14
 - supported versions 21-3
 - support for 1-4
 - VLAN configuration 21-8
 - IGMP throttling
 - configuring 21-27
 - default configuration 21-24
 - described 21-24
 - displaying action 21-28
 - Immediate Leave, IGMP 21-5
 - enabling 36-9
 - inaccessible authentication bypass 10-24
 - support for multiauth ports 10-25
 - initial configuration
 - defaults 1-16
 - Express Setup 1-2
 - interface
 - number 12-16
 - range macros 12-19
 - interface command 12-15 to ??, 12-15 to 12-16
 - interface configuration mode 2-2
 - interfaces
 - auto-MDIX, configuring 12-29
 - configuration guidelines
 - duplex and speed 12-26
 - configuring
 - procedure 12-16
 - counters, clearing 12-38
 - default configuration 12-23
 - described 12-34
 - descriptive name, adding 12-34
 - displaying information about 12-37
 - flow control 12-28
 - management 1-5
 - monitoring 12-37
 - naming 12-34
 - physical, identifying 12-15

- range of [12-17](#)
 - restarting [12-38](#)
 - shutting down [12-38](#)
 - speed and duplex, configuring [12-27](#)
 - status [12-37](#)
 - supported [12-15](#)
 - types of [12-1](#)
- interfaces range macro command [12-19](#)
- interface types [12-15](#)
- Internet Protocol version 6
- See IPv6
- inter-VLAN routing [34-1](#)
- Intrusion Detection System
- See IDS appliances
- inventory management TLV [26-3, 26-7](#)
- IP ACLs
- for QoS classification [33-8](#)
 - implicit deny [31-9, 31-12](#)
 - implicit masks [31-9](#)
 - named [31-13](#)
 - undefined [31-19](#)
- IP addresses
- 128-bit [35-2](#)
 - candidate or member [6-4, 6-12](#)
 - classes of [34-4](#)
 - cluster access [6-2](#)
 - command switch [6-3, 6-10, 6-12](#)
 - discovering [5-24](#)
 - for IP routing [34-4](#)
 - IPv6 [35-2](#)
 - redundant clusters [6-10](#)
 - standby command switch [6-10, 6-12](#)
 - See also IP information
- ip igmp profile command [21-24](#)
- IP information
- assigned
 - manually [3-14](#)
 - through DHCP-based autoconfiguration [3-3](#)
 - default configuration [3-3](#)
- IP phones
- and QoS [15-1](#)
 - automatic classification and queueing [33-21](#)
 - configuring [15-4](#)
 - ensuring port security with QoS [33-45](#)
 - trusted boundary for QoS [33-45](#)
- IP Port Security for Static Hosts
- on a Layer 2 access port [20-17](#)
- IP precedence [33-2](#)
- IP-precedence-to-DSCP map for QoS [33-63](#)
- IP protocols in ACLs [31-10](#)
- IP routing
- disabling [34-4](#)
 - enabling [34-4](#)
- IP Service Level Agreements
- See IP SLAs
- IP service levels, analyzing [32-1](#)
- IP SLAs
- benefits [32-2](#)
 - configuration guidelines [32-5](#)
 - Control Protocol [32-4](#)
 - default configuration [32-5](#)
 - definition [32-1](#)
 - measuring network performance [32-3](#)
 - monitoring [32-6](#)
 - operation [32-3](#)
 - responder
 - described [32-4](#)
 - enabling [32-6](#)
 - response time [32-4](#)
 - SNMP support [32-2](#)
 - supported metrics [32-2](#)
- IP source guard
- and 802.1x [20-16](#)
 - and DHCP snooping [20-13](#)
 - and EtherChannels [20-16](#)
 - and port security [20-16](#)
 - and private VLANs [20-16](#)
 - and routed ports [20-15](#)

- and TCAM entries 20-16
- and trunk interfaces 20-15
- and VRF 20-16
- binding configuration
 - automatic 20-13
 - manual 20-13
- binding table 20-13
- configuration guidelines 20-15
- default configuration 20-15
- described 20-13
- disabling 20-17
- displaying
 - active IP or MAC bindings 20-21
 - bindings 20-21
 - configuration 20-21
- enabling 20-16, 20-17
- filtering
 - source IP address 20-14
 - source IP and MAC address 20-14
- on provisioned switches 20-16
- source IP address filtering 20-14
- source IP and MAC address filtering 20-14
- static bindings
 - adding 20-16, 20-17
 - deleting 20-17
- static hosts 20-17

IP traceroute

- executing 38-18
- overview 38-17

IP unicast routing

- assigning IP addresses to Layer 3 interfaces 34-4
- configuring static routes 34-5
- disabling 34-4
- enabling 34-4
- inter-VLAN 34-1
- IP addressing
 - classes 34-4
 - configuring 34-4
- steps to configure 34-3

- subnet mask 34-4
- with SVIs 34-3

IPv4 ACLs

- applying to interfaces 31-17
- extended, creating 31-9
- named 31-13
- standard, creating 31-8

IPv4 and IPv6

- dual protocol stacks 35-4

IPv6

- addresses 35-2
- address formats 35-2
- and switch stacks 35-6
- applications 35-4
- assigning address 35-7
- autoconfiguration 35-4
- configuring static routes 35-10
- default configuration 35-6
- defined 35-1
- forwarding 35-7
- ICMP 35-3
- monitoring 35-11
- neighbor discovery 35-3
- SDM templates 36-1
- stack master functions 35-6
- Stateless Autoconfiguration 35-4
- supported features 35-2

J

- join messages, IGMP 21-3

L

LACP

- See EtherChannel

Layer 2 frames, classification with CoS 33-2

Layer 2 interfaces, default configuration 12-23

Layer 2 traceroute

- and ARP [38-16](#)
- and CDP [38-16](#)
- broadcast traffic [38-15](#)
- described [38-15](#)
- IP addresses and subnets [38-16](#)
- MAC addresses and VLANs [38-16](#)
- multicast traffic [38-16](#)
- multiple devices on a port [38-16](#)
- unicast traffic [38-15](#)
- usage guidelines [38-16](#)

Layer 3 features [1-15](#)

Layer 3 interfaces

- assigning IP addresses to [34-4](#)
- assigning IPv6 addresses to [35-7](#)
- changing from Layer 2 mode [34-4](#)

Layer 3 packets, classification methods [33-2](#)LDAP [4-2](#)Leaking IGMP Reports [19-4](#)

LEDs, switch

- See hardware installation guide

lightweight directory access protocol

- See LDAP

line configuration mode [2-2](#)

Link Aggregation Control Protocol

- See EtherChannel

link failure, detecting unidirectional [17-8](#)

Link Layer Discovery Protocol

- See CDP

link local unicast addresses [35-3](#)

link redundancy

- See Flex Links

links, unidirectional [24-1](#)

link-state tracking

- configuring [37-23](#)
- described [37-20](#)

LLDP

- configuring [26-5](#)
 - characteristics [26-6](#)
 - default configuration [26-5](#)
- enabling [26-6](#)
- monitoring and maintaining [26-11](#)
- overview [26-1](#)
- supported TLVs [26-2](#)
- switch stack considerations [26-2](#)
- transmission timer and holdtime, setting [26-6](#)

LLDP-MED

- configuring
 - procedures [26-5](#)
 - TLVs [26-7](#)
- monitoring and maintaining [26-11](#)
- overview [26-1, 26-2](#)
- supported TLVs [26-2](#)

LLDP Media Endpoint Discovery

- See LLDP-MED

local SPAN [27-2](#)location TLV [26-3, 26-7](#)

login authentication

- with RADIUS [9-30](#)
- with TACACS+ [9-14](#)

login banners [5-11](#)

log messages

- See system message logging

Long-Reach Ethernet (LRE) technology [1-20](#)

loop guard

- described [18-11](#)
- enabling [18-18](#)
- support for [1-9](#)

LRE profiles, considerations in switch clusters [6-15](#)

M

MAB

- See MAC authentication bypass

- MAB inactivity timer
 - default setting [10-35](#)
 - range [10-37](#)
- MAC/PHY configuration status TLV [26-2](#)
- MAC addresses
 - aging time [5-15](#)
 - and VLAN association [5-14](#)
 - building the address table [5-14](#)
 - default configuration [5-15](#)
 - disabling learning on a VLAN [5-23](#)
 - discovering [5-24](#)
 - displaying [5-24](#)
 - displaying in the IP source binding table [20-20](#)
 - dynamic
 - learning [5-14](#)
 - removing [5-16](#)
 - in ACLs [31-22](#)
 - static
 - adding [5-21](#)
 - allowing [5-22, 5-23](#)
 - characteristics of [5-20](#)
 - dropping [5-22](#)
 - removing [5-21](#)
- MAC address learning [1-6](#)
- MAC address learning, disabling on a VLAN [5-23](#)
- MAC address notification, support for [1-15](#)
- MAC address-table move update
 - configuration guidelines [19-8](#)
 - configuring [19-12](#)
 - default configuration [19-8](#)
 - description [19-6](#)
 - monitoring [19-14](#)
- MAC address-to-VLAN mapping [13-23](#)
- MAC authentication bypass [10-37](#)
 - configuring [10-56](#)
 - overview [10-17](#)
- MAC extended access lists
 - applying to Layer 2 interfaces [31-23](#)
 - configuring for QoS [33-51](#)
 - creating [31-22](#)
 - defined [31-22](#)
 - for QoS classification [33-5](#)
- magic packet [10-27](#)
- manageability features [1-6](#)
- management access
 - in-band
 - browser session [1-7](#)
 - CLI session [1-7](#)
 - device manager [1-7](#)
 - SNMP [1-7](#)
 - out-of-band console port connection [1-7](#)
- management address TLV [26-2](#)
- management options
 - CLI [2-1](#)
 - clustering [1-3](#)
 - CNS [4-1](#)
 - Network Assistant [1-2](#)
 - overview [1-5](#)
- management VLAN
 - considerations in switch clusters [6-7](#)
 - discovery through different management VLANs [6-7](#)
- mapping tables for QoS
 - configuring
 - CoS-to-DSCP [33-62](#)
 - DSCP [33-62](#)
 - DSCP-to-CoS [33-65](#)
 - DSCP-to-DSCP-mutation [33-66](#)
 - IP-precedence-to-DSCP [33-63](#)
 - policed-DSCP [33-64](#)
 - described [33-11](#)
- marking
 - action with aggregate policers [33-59](#)
 - described [33-4, 33-9](#)
- matching, IPv4 ACLs [31-7](#)

- maximum aging time
 - MSTP 17-24
 - STP 16-23
- maximum hop count, MSTP 17-25
- maximum number of allowed devices, port-based authentication 10-38
- MDA
 - configuration guidelines 10-12 to 10-13
 - described 1-11, 10-12
 - exceptions with authentication process 10-5
- membership mode, VLAN port 13-3
- member switch
 - automatic discovery 6-5
 - defined 6-2
 - managing 6-15
 - passwords 6-12
 - recovering from lost connectivity 38-12
 - requirements 6-4
 - See also candidate switch, cluster standby group, and standby command switch
- memory consistency check errors
 - example 38-26
- memory consistency check routines 1-5, 38-26
- memory consistency integrity 1-5, 38-26
- messages, to users through banners 5-11
- MIBs
 - overview 30-1
 - SNMP interaction with 30-5
- mirroring traffic for analysis 27-1
- mismatches, autonegotiation 38-12
- module number 12-16
- monitoring
 - access groups 31-24
 - cables for unidirectional links 24-1
 - CDP 25-5
 - features 1-15
 - Flex Links 19-14
 - IGMP
 - filters 21-28
 - snooping 21-15, 36-11
 - interfaces 12-37
 - IP SLAs operations 32-6
 - IPv4 ACL configuration 31-24
 - IPv6 35-11
 - MAC address-table move update 19-14
 - multicast router interfaces 21-16, 36-12
 - MVR 21-22
 - network traffic for analysis with probe 27-2
 - port
 - blocking 23-20
 - protection 23-20
 - SFP status 12-37, 38-14
 - speed and duplex mode 12-27
 - traffic flowing among switches 28-1
 - traffic suppression 23-20
 - VLANs 13-13
 - VMPS 13-27
 - VTP 14-17
- mrouter Port 19-3
- mrouter port 19-5
- MSTP
 - boundary ports
 - configuration guidelines 17-15
 - described 17-6
 - BPDU filtering
 - described 18-3
 - enabling 18-15
 - BPDU guard
 - described 18-2
 - enabling 18-14
 - CIST, described 17-3
 - CIST regional root 17-3
 - CIST root 17-5
 - configuration guidelines 17-15, 18-12
 - configuring
 - forward-delay time 17-24
 - hello time 17-23
 - link type for rapid convergence 17-25

- maximum aging time [17-24](#)
 - maximum hop count [17-25](#)
 - MST region [17-16](#)
 - neighbor type [17-26](#)
 - path cost [17-21](#)
 - port priority [17-20](#)
 - root switch [17-18](#)
 - secondary root switch [17-19](#)
 - switch priority [17-22](#)
- CST
 - defined [17-3](#)
 - operations between regions [17-4](#)
- default configuration [17-14](#)
- default optional feature configuration [18-12](#)
- displaying status [17-27](#)
- enabling the mode [17-16](#)
- EtherChannel guard
 - described [18-10](#)
 - enabling [18-17](#)
- extended system ID
 - effects on root switch [17-18](#)
 - effects on secondary root switch [17-19](#)
 - unexpected behavior [17-18](#)
- IEEE 802.1s
 - implementation [17-6](#)
 - port role naming change [17-7](#)
 - terminology [17-5](#)
- instances supported [16-10](#)
- interface state, blocking to forwarding [18-2](#)
- interoperability and compatibility among modes [16-11](#)
- interoperability with IEEE 802.1D
 - described [17-9](#)
 - restarting migration process [17-26](#)
- IST
 - defined [17-3](#)
 - master [17-3](#)
 - operations within a region [17-3](#)
- loop guard
 - described [18-11](#)
 - enabling [18-18](#)
- mapping VLANs to MST instance [17-16](#)
- MST region
 - CIST [17-3](#)
 - configuring [17-16](#)
 - described [17-2](#)
 - hop-count mechanism [17-5](#)
 - IST [17-3](#)
 - supported spanning-tree instances [17-2](#)
- optional features supported [1-9](#)
- overview [17-2](#)
- Port Fast
 - described [18-2](#)
 - enabling [18-12](#)
- preventing root switch selection [18-10](#)
- root guard
 - described [18-10](#)
 - enabling [18-18](#)
- root switch
 - configuring [17-18](#)
 - effects of extended system ID [17-18](#)
 - unexpected behavior [17-18](#)
- shutdown Port Fast-enabled port [18-2](#)
- stack changes, effects of [17-8](#)
- status, displaying [17-27](#)
- multiauth
 - support for inaccessible authentication bypass [10-25](#)
- multiauth mode
 - See multiple-authentication mode
- multicast groups
 - Immediate Leave [21-5](#)
 - joining [21-3](#)
 - leaving [21-5](#)
 - static joins [21-10, 36-8](#)
- multicast router interfaces, monitoring [21-16, 36-12](#)
- multicast router ports, adding [21-9, 36-8](#)
- multicast storm [23-1](#)
- multicast storm-control command [23-4](#)

multicast television application [21-18](#)

multicast VLAN [21-17](#)

Multicast VLAN Registration

See MVR

multidomain authentication

See MDA

multiple authentication [10-13](#)

multiple authentication mode

configuring [10-44](#)

MVR

and address aliasing [21-20](#)

and IGMPv3 [21-20](#)

configuration guidelines [21-20](#)

configuring interfaces [21-21](#)

default configuration [21-19](#)

described [21-17](#)

example application [21-18](#)

modes [21-21](#)

monitoring [21-22](#)

multicast television application [21-18](#)

setting global parameters [21-20](#)

support for [1-5](#)

N

NAC

critical authentication [10-24, 10-54](#)

IEEE 802.1x authentication using a RADIUS server [10-58](#)

IEEE 802.1x validation using RADIUS server [10-58](#)

inaccessible authentication bypass [10-54](#)

Layer 2 IEEE 802.1x validation [1-12, 10-29, 10-58](#)

named IPv4 ACLs [31-13](#)

NameSpace Mapper

See NSM

native VLAN

configuring [13-18](#)

default [13-18](#)

NEAT

configuring [10-59](#)

overview [10-31](#)

neighbor discovery, IPv6 [35-3](#)

Network Admission Control

See NAC

Network Assistant

benefits [1-2](#)

described [1-5](#)

downloading image files [1-2](#)

guide mode [1-2](#)

management options [1-2](#)

managing switch stacks [7-2, 7-14](#)

upgrading a switch [A-24](#)

wizards [1-2](#)

network configuration examples

increasing network performance [1-19](#)

long-distance, high-bandwidth transport [1-23](#)

providing network services [1-19](#)

server aggregation and Linux server cluster [1-21](#)

small to medium-sized network [1-22](#)

network design

performance [1-19](#)

services [1-19](#)

Network Edge Access Topology

See NEAT

network management

CDP [25-1](#)

RMON [28-1](#)

SNMP [30-1](#)

network performance, measuring with IP SLAs [32-3](#)

network policy TLV [26-2, 26-8](#)

Network Time Protocol

See NTP

no commands [2-4](#)

nonhierarchical policy maps

described [33-10](#)

non-IP traffic filtering [31-22](#)

nontrunking mode [13-14](#)

normal-range VLANs [13-4](#)
 configuration guidelines [13-6](#)
 configuring [13-4](#)
 defined [13-1](#)

NSM [4-3](#)

NTP

 associations
 defined [5-3](#)
 overview [5-3](#)
 stratum [5-3](#)
 support for [1-7](#)
 time
 services [5-3](#)
 synchronizing [5-3](#)

O

OBFL

 configuring [38-25](#)
 described [38-24](#)
 displaying [38-25](#)

offline configuration for switch stacks [7-7](#)

off mode, VTP [14-4](#)

on-board failure logging

 See OBFL

online diagnostics

 overview [39-1](#)
 running tests [39-3](#)
 understanding [39-1](#)

openlx

 configuring [10-64](#)

openlx authentication

 overview [10-30](#)

optimizing system resources [8-1](#)

options, management [1-5](#)

out-of-profile markdown [1-14](#)

P

packet modification, with QoS [33-20](#)

PAgP

 See EtherChannel

passwords

 default configuration [9-2](#)
 disabling recovery of [9-5](#)
 encrypting [9-3](#)
 for security [1-10](#)
 in clusters [6-13](#)
 overview [9-1](#)
 recovery of [38-3](#)
 setting
 enable [9-3](#)
 enable secret [9-3](#)
 Telnet [9-6](#)
 with usernames [9-6](#)
 VTP domain [14-10](#)

path cost

 MSTP [17-21](#)
 STP [16-20](#)

PC (passive command switch) [6-9](#)

performance, network design [1-19](#)

performance features [1-4](#)

persistent self-signed certificate [9-47](#)

per-user ACLs and Filter-Ids [10-8](#)

per-VLAN spanning-tree plus

 See PVST+

physical ports [12-2](#)

PIM-DVMRP, as snooping method [21-8](#)

ping

 character output description [38-15](#)
 executing [38-14](#)
 overview [38-14](#)

PoE

 auto mode [12-7](#)

 CDP with power consumption, described [12-5](#)

 CDP with power negotiation, described [12-5](#)

- Cisco intelligent power management [12-5](#)
 - configuring [12-30](#)
 - cutoff power
 - determining [12-8](#)
 - cutoff-power
 - support for [12-8](#)
 - devices supported [12-5](#)
 - high-power devices operating in low-power mode [12-5](#)
 - IEEE power classification levels [12-6](#)
 - monitoring [12-8](#)
 - monitoring power [12-33](#)
 - policing power consumption [12-33](#)
 - policing power usage [12-8](#)
 - power budgeting [12-31](#)
 - power consumption [12-9, 12-31](#)
 - powered-device detection and initial power allocation [12-6](#)
 - power management modes [12-7](#)
 - power monitoring [12-8](#)
 - power negotiation extensions to CDP [12-5](#)
 - power sensing [12-8](#)
 - standards supported [12-5](#)
 - static mode [12-7](#)
 - total available power [12-9](#)
 - troubleshooting [38-12](#)
- PoE+ [1-15, 12-5, 12-6, 12-30](#)
- policed-DSCP map for QoS [33-64](#)
- policers
 - configuring
 - for each matched traffic class [33-54](#)
 - for more than one traffic class [33-59](#)
 - described [33-4](#)
 - displaying [33-80](#)
 - number of [33-41](#)
 - types of [33-10](#)
- policing
 - described [33-4](#)
 - token-bucket algorithm [33-10](#)
- policy maps for QoS
 - characteristics of [33-54](#)
 - described [33-8](#)
 - displaying [33-81](#)
 - nonhierarchical on physical ports
 - described [33-10](#)
- port ACLs
 - defined [31-2](#)
 - types of [31-3](#)
- Port Aggregation Protocol
 - See EtherChannel
- port-based authentication
 - accounting [10-16](#)
 - authentication server
 - defined [10-3, 11-2](#)
 - RADIUS server [10-3](#)
 - client, defined [10-3, 11-2](#)
 - configuration guidelines [10-36, 11-9](#)
 - configuring
 - 802.1x authentication [10-41](#)
 - guest VLAN [10-51](#)
 - host mode [10-44](#)
 - inaccessible authentication bypass [10-54](#)
 - manual re-authentication of a client [10-46](#)
 - periodic re-authentication [10-45](#)
 - quiet period [10-46](#)
 - RADIUS server [10-44, 11-13](#)
 - RADIUS server parameters on the switch [10-43, 11-11](#)
 - restricted VLAN [10-52](#)
 - switch-to-client frame-retransmission number [10-47, 10-48](#)
 - switch-to-client retransmission time [10-47](#)
 - violation modes [10-40 to 10-41](#)
- default configuration [10-35, 11-9](#)
- described [10-1](#)
- device roles [10-3, 11-2](#)
- displaying statistics [10-65, 11-17](#)
- downloadable ACLs and redirect URLs

- configuring [10-60 to 10-62, ?? to 10-63](#)
 - overview [10-20 to 10-22](#)
- EAPOL-start frame [10-5](#)
- EAP-request/identity frame [10-5](#)
- EAP-response/identity frame [10-5](#)
- enabling
 - 802.1X authentication [11-11](#)
- encapsulation [10-3](#)
- flexible authentication ordering
 - configuring [10-63](#)
 - overview [10-30](#)
- guest VLAN
 - configuration guidelines [10-23, 10-24](#)
 - described [10-22](#)
- host mode [10-12](#)
- inaccessible authentication bypass
 - configuring [10-54](#)
 - described [10-24](#)
 - guidelines [10-37](#)
- initiation and message exchange [10-5](#)
- magic packet [10-27](#)
- maximum number of allowed devices per port [10-38](#)
- method lists [10-41](#)
- multiple authentication [10-13](#)
- per-user ACLs
 - configuration tasks [10-20](#)
 - described [10-19](#)
 - RADIUS server attributes [10-19](#)
- ports
 - authorization state and dot1x port-control command [10-10](#)
 - authorized and unauthorized [10-10](#)
 - voice VLAN [10-26](#)
- port security
 - described [10-27](#)
- readiness check
 - configuring [10-38](#)
 - described [10-17, 10-38](#)
- resetting to default values [10-65](#)
- stack changes, effects of [10-11](#)
- statistics, displaying [10-65](#)
- switch
 - as proxy [10-3, 11-2](#)
 - RADIUS client [10-3](#)
- switch supplicant
 - configuring [10-59](#)
 - overview [10-31](#)
- user distribution
 - guidelines [10-29](#)
 - overview [10-29](#)
- VLAN assignment
 - AAA authorization [10-41](#)
 - characteristics [10-18](#)
 - configuration tasks [10-18](#)
 - described [10-17](#)
- voice aware 802.1x security
 - configuring [10-39](#)
 - described [10-31, 10-39](#)
- voice VLAN
 - described [10-26](#)
 - PVID [10-26](#)
 - VVID [10-26](#)
- wake-on-LAN, described [10-27](#)
- with ACLs and RADIUS Filter-Id attribute [10-32](#)
- port-based authentication methods, supported [10-7](#)
- port blocking [1-4, 23-7](#)
- port-channel
 - See EtherChannel
- port description TLV [26-2](#)
- Port Fast
 - described [18-2](#)
 - enabling [18-12](#)
 - mode, spanning tree [13-24](#)
 - support for [1-9](#)
- port membership modes, VLAN [13-3](#)
- port priority
 - MSTP [17-20](#)
 - STP [16-18](#)

- ports
 - access [12-3](#)
 - blocking [23-7](#)
 - dual-purpose uplink [12-4](#)
 - dynamic access [13-4](#)
 - protected [23-6](#)
 - secure [23-9](#)
 - static-access [13-3, 13-9](#)
 - switch [12-2](#)
 - trunks [13-3, 13-13](#)
 - VLAN assignments [13-9](#)
 - port security
 - aging [23-17](#)
 - and QoS trusted boundary [33-45](#)
 - and stacking [23-18](#)
 - configuring [23-12](#)
 - default configuration [23-11](#)
 - described [23-8](#)
 - displaying [23-20](#)
 - on trunk ports [23-14](#)
 - sticky learning [23-9](#)
 - violations [23-10](#)
 - with other features [23-11](#)
 - port-shutdown response, VMPS [13-23](#)
 - port VLAN ID TLV [26-2](#)
 - power management TLV [26-3, 26-8](#)
 - Power over Ethernet
 - See PoE
 - preemption, default configuration [19-7](#)
 - preemption delay, default configuration [19-8](#)
 - preferential treatment of traffic
 - See QoS
 - preventing unauthorized access [9-1](#)
 - primary links [19-2](#)
 - priority
 - overriding CoS [15-6](#)
 - trusting CoS [15-6](#)
 - private VLAN edge ports
 - See protected ports
 - privileged EXEC mode [2-2](#)
 - privilege levels
 - changing the default for lines [9-9](#)
 - command switch [6-16](#)
 - exiting [9-9](#)
 - logging into [9-9](#)
 - mapping on member switches [6-16](#)
 - overview [9-2, 9-7](#)
 - setting a command with [9-8](#)
 - protected ports [1-10, 23-6](#)
 - protocol storm protection [23-18](#)
 - provisioned switches and IP source guard [20-16](#)
 - provisioning new members for a switch stack [7-7](#)
 - proxy reports [19-4](#)
 - pruning, VTP
 - disabling
 - in VTP domain [14-16](#)
 - on a port [13-18](#)
 - enabling
 - in VTP domain [14-15](#)
 - on a port [13-18](#)
 - examples [14-7](#)
 - overview [14-6](#)
 - pruning-eligible list
 - changing [13-18](#)
 - for VTP pruning [14-6](#)
 - VLANs [14-16](#)
 - PVST+
 - described [16-10](#)
 - IEEE 802.1Q trunking interoperability [16-11](#)
 - instances supported [16-10](#)
-
- ## Q
- QoS
 - and MQC commands [33-1](#)
 - auto-QoS
 - categorizing traffic [33-22](#)
 - configuration and defaults display [33-37](#)

- configuration guidelines [33-34](#)
- described [33-21](#)
- disabling [33-36](#)
- displaying generated commands [33-36](#)
- displaying the initial configuration [33-37](#)
- effects on running configuration [33-34](#)
- list of generated commands [33-25, 33-29](#)
- basic model [33-4](#)
- classification
 - class maps, described [33-8](#)
 - defined [33-4](#)
 - DSCP transparency, described [33-46](#)
 - flowchart [33-7](#)
 - forwarding treatment [33-3](#)
 - in frames and packets [33-3](#)
 - IP ACLs, described [33-6, 33-8](#)
 - MAC ACLs, described [33-5, 33-8](#)
 - options for IP traffic [33-6](#)
 - options for non-IP traffic [33-5](#)
 - policy maps, described [33-8](#)
 - trust DSCP, described [33-5](#)
 - trusted CoS, described [33-5](#)
 - trust IP precedence, described [33-5](#)
- class maps
 - configuring [33-52](#)
 - displaying [33-80](#)
- configuration guidelines
 - auto-QoS [33-34](#)
 - standard QoS [33-40](#)
- configuring
 - aggregate policers [33-59](#)
 - auto-QoS [33-21](#)
 - default port CoS value [33-44](#)
 - DSCP maps [33-62](#)
 - DSCP transparency [33-46](#)
 - DSCP trust states bordering another domain [33-47](#)
 - egress queue characteristics [33-73](#)
 - ingress queue characteristics [33-68](#)
 - IP extended ACLs [33-50](#)
 - IP standard ACLs [33-49](#)
 - MAC ACLs [33-51](#)
 - port trust states within the domain [33-42](#)
 - trusted boundary [33-45](#)
- default auto configuration [33-22](#)
- default standard configuration [33-38](#)
- displaying statistics [33-80](#)
- DSCP transparency [33-46](#)
- egress queues
 - allocating buffer space [33-73](#)
 - buffer allocation scheme, described [33-18](#)
 - configuring shaped weights for SRR [33-77](#)
 - configuring shared weights for SRR [33-78](#)
 - described [33-4](#)
 - displaying the threshold map [33-76](#)
 - flowchart [33-18](#)
 - mapping DSCP or CoS values [33-75](#)
 - scheduling, described [33-4](#)
 - setting WTD thresholds [33-73](#)
 - WTD, described [33-19](#)
- enabling globally [33-42](#)
- flowcharts
 - classification [33-7](#)
 - egress queueing and scheduling [33-18](#)
 - ingress queueing and scheduling [33-15](#)
 - policing and marking [33-11](#)
- implicit deny [33-8](#)
- ingress queues
 - allocating bandwidth [33-71](#)
 - allocating buffer space [33-70](#)
 - buffer and bandwidth allocation, described [33-16](#)
 - configuring shared weights for SRR [33-71](#)
 - configuring the priority queue [33-72](#)
 - described [33-4](#)
 - displaying the threshold map [33-69](#)
 - flowchart [33-15](#)
 - mapping DSCP or CoS values [33-68](#)

- priority queue, described 33-17
 - scheduling, described 33-4
 - setting WTD thresholds 33-68
 - WTD, described 33-16
 - IP phones
 - automatic classification and queueing 33-21
 - detection and trusted settings 33-21, 33-45
 - limiting bandwidth on egress interface 33-79
 - mapping tables
 - CoS-to-DSCP 33-62
 - displaying 33-81
 - DSCP-to-CoS 33-65
 - DSCP-to-DSCP-mutation 33-66
 - IP-precedence-to-DSCP 33-63
 - policed-DSCP 33-64
 - types of 33-11
 - marked-down actions 33-57
 - marking, described 33-4, 33-9
 - overview 33-2
 - packet modification 33-20
 - policers
 - configuring 33-57, 33-60
 - described 33-9
 - displaying 33-80
 - number of 33-41
 - types of 33-10
 - policies, attaching to an interface 33-9
 - policing
 - described 33-4, 33-9
 - token bucket algorithm 33-10
 - policy maps
 - characteristics of 33-54
 - displaying 33-81
 - nonhierarchical on physical ports 33-54
 - QoS label, defined 33-4
 - queues
 - configuring egress characteristics 33-73
 - configuring ingress characteristics 33-68
 - high priority (expedite) 33-20, 33-79
 - location of 33-12
 - SRR, described 33-14
 - WTD, described 33-13
 - rewrites 33-20
 - support for 1-13
 - trust states
 - bordering another domain 33-47
 - described 33-5
 - trusted device 33-45
 - within the domain 33-42
 - quality of service
 - See QoS
 - queries, IGMP 21-4
 - query solicitation, IGMP 21-13
-
- ## R
- RADIUS
 - attributes
 - vendor-proprietary 9-38
 - vendor-specific 9-36
 - configuring
 - accounting 9-35
 - authentication 9-30
 - authorization 9-34
 - communication, global 9-28, 9-36
 - communication, per-server 9-27, 9-28
 - multiple UDP ports 9-27
 - default configuration 9-27
 - defining AAA server groups 9-32
 - displaying the configuration 9-40
 - identifying the server 9-27
 - in clusters 6-15
 - limiting the services to the user 9-34
 - method list, defined 9-26
 - operation of 9-19
 - overview 9-18
 - server load balancing 9-40
 - suggested network environments 9-18

- support for [1-12](#)
 - tracking services accessed by user [9-35](#)
- RADIUS Change of Authorization [9-20](#)
- range
 - macro [12-19](#)
 - of interfaces [12-17](#)
- rapid convergence [17-10](#)
- rapid per-VLAN spanning-tree plus
 - See rapid PVST+
- rapid PVST+
 - described [16-10](#)
 - IEEE 802.1Q trunking interoperability [16-11](#)
 - instances supported [16-10](#)
- Rapid Spanning Tree Protocol
 - See RSTP
- rcommand command [6-15](#)
- RCP
 - configuration files
 - downloading [A-17](#)
 - overview [A-16](#)
 - preparing the server [A-16](#)
 - uploading [A-18](#)
 - image files
 - deleting old image [A-37](#)
 - downloading [A-35](#)
 - preparing the server [A-34](#)
 - uploading [A-37](#)
- readiness check
 - port-based authentication
 - configuring [10-38](#)
 - described [10-17, 10-38](#)
- reconfirmation interval, VMPS, changing [13-26](#)
- reconfirming dynamic VLAN membership [13-26](#)
- recovery procedures [38-1](#)
- redirect URL [10-20, 10-21, 10-60](#)
- redundancy
 - EtherChannel [37-3](#)
 - STP
 - backbone [16-8](#)
 - multidrop backbone [18-5](#)
 - path cost [13-21](#)
 - port priority [13-19](#)
- redundant links and UplinkFast [18-16](#)
- reloading software [3-22](#)
- Remote Authentication Dial-In User Service
 - See RADIUS
- Remote Copy Protocol
 - See RCP
- Remote Network Monitoring
 - See RMON
- Remote SPAN
 - See RSPAN
- remote SPAN [27-3](#)
- report suppression, IGMP
 - described [21-6](#)
 - disabling [21-15, 36-11](#)
- resequencing ACL entries [31-13](#)
- reserved addresses in DHCP pools [20-22](#)
- resetting a UDLD-shutdown interface [24-6](#)
- responder, IP SLAs
 - described [32-4](#)
 - enabling [32-6](#)
- response time, measuring with IP SLAs [32-4](#)
- restricted VLAN
 - configuring [10-52](#)
 - described [10-23](#)
 - using with IEEE 802.1x [10-23](#)
- restricting access
 - overview [9-1](#)
 - passwords and privilege levels [9-2](#)
 - RADIUS [9-17](#)
 - TACACS+ [9-10](#)
- retry count, VMPS, changing [13-27](#)

- RFC
 - 1112, IP multicast and IGMP [21-2](#)
 - 1157, SNMPv1 [30-2](#)
 - 1166, IP addresses [34-4](#)
 - 1305, NTP [5-3](#)
 - 1757, RMON [28-2](#)
 - 1901, SNMPv2C [30-2](#)
 - 1902 to 1907, SNMPv2 [30-2](#)
 - 2236, IP multicast and IGMP [21-2](#)
 - 2273-2275, SNMPv3 [30-2](#)
- RFC 5176 Compliance [9-21](#)
- RMON
 - default configuration [28-3](#)
 - displaying status [28-6](#)
 - enabling alarms and events [28-3](#)
 - groups supported [28-2](#)
 - overview [28-1](#)
 - statistics
 - collecting group Ethernet [28-5](#)
 - collecting group history [28-5](#)
 - support for [1-16](#)
- root guard
 - described [18-10](#)
 - enabling [18-18](#)
 - support for [1-9](#)
- root switch
 - MSTP [17-18](#)
 - STP [16-16](#)
- router ACLs
 - defined [31-2](#)
 - types of [31-4](#)
- RSPAN
 - and stack changes [27-10](#)
 - characteristics [27-8](#)
 - configuration guidelines [27-17](#)
 - default configuration [27-10](#)
 - defined [27-3](#)
 - destination ports [27-7](#)
 - displaying status [27-23](#)
 - in a switch stack [27-2](#)
 - interaction with other features [27-9](#)
 - monitored ports [27-6](#)
 - monitoring ports [27-7](#)
 - overview [1-16, 27-1](#)
 - received traffic [27-5](#)
 - sessions
 - creating [27-18](#)
 - defined [27-4](#)
 - limiting source traffic to specific VLANs [27-22](#)
 - specifying monitored ports [27-18](#)
 - with ingress traffic enabled [27-21](#)
 - source ports [27-6](#)
 - transmitted traffic [27-6](#)
 - VLAN-based [27-7](#)
- RSTP
 - active topology [17-10](#)
- BPDU
 - format [17-12](#)
 - processing [17-13](#)
- designated port, defined [17-9](#)
- designated switch, defined [17-9](#)
- interoperability with IEEE 802.1D
 - described [17-9](#)
 - restarting migration process [17-26](#)
 - topology changes [17-13](#)
- overview [17-9](#)
- port roles
 - described [17-9](#)
 - synchronized [17-11](#)
- proposal-agreement handshake process [17-10](#)
- rapid convergence
 - cross-stack rapid convergence [17-11](#)
 - described [17-10](#)
 - edge ports and Port Fast [17-10](#)
 - point-to-point links [17-10, 17-25](#)
 - root ports [17-10](#)
- root port, defined [17-9](#)
- See also MSTP

running configuration
 replacing [A-19, A-20](#)
 rolling back [A-19, A-21](#)
 running configuration, saving [3-15](#)

S

SC (standby command switch) [6-9](#)

scheduled reloads [3-22](#)

SCP

and SSH [9-52](#)

configuring [9-53](#)

SDM

templates

 configuring [8-4](#)

 number of [8-1](#)

SDM template

 configuration guidelines [8-3](#)

 configuring [8-3](#)

 types of [8-1](#)

Secure Copy Protocol

secure HTTP client

 configuring [9-51](#)

 displaying [9-52](#)

secure HTTP server

 configuring [9-50](#)

 displaying [9-52](#)

secure MAC addresses

 and switch stacks [23-18](#)

 deleting [23-16](#)

 maximum number of [23-10](#)

 types of [23-9](#)

secure ports

 and switch stacks [23-18](#)

secure ports, configuring [23-9](#)

secure remote connections [9-42](#)

Secure Shell

 See SSH

Secure Socket Layer

 See SSL

security, port [23-8](#)

security features [1-10](#)

See SCP

sequence numbers in log messages [29-8](#)

server mode, VTP [14-3](#)

service-provider network, MSTP and RSTP [17-1](#)

set-request operation [30-5](#)

setup program

 failed command switch replacement [38-10](#)

 replacing failed command switch [38-9](#)

severity levels, defining in system messages [29-9](#)

SFPs

 monitoring status of [12-37, 38-14](#)

 security and identification [38-13](#)

 status, displaying [38-14](#)

shaped round robin

 See SRR

show access-lists hw-summary command [31-19](#)

show and more command output, filtering [2-9](#)

show cdp traffic command [25-5](#)

show cluster members command [6-15](#)

show configuration command [12-34](#)

show forward command [38-21](#)

show interfaces command [12-27, 12-34](#)

show interfaces switchport [19-4](#)

show lldp traffic command [26-12](#)

show platform forward command [38-21](#)

show platform team command [38-26](#)

show running-config command

 displaying ACLs [31-17, 31-18](#)

 interface description in [12-34](#)

shutdown command on interfaces [12-38](#)

Simple Network Management Protocol

 See SNMP

small form-factor pluggable modules

 See SFPs

small-frame arrival rate, configuring [23-5](#)

SNAP 25-1

SNMP

accessing MIB variables with 30-5

agent

described 30-4

disabling 30-8

and IP SLAs 32-2

authentication level 30-11

community strings

configuring 30-8

for cluster switches 30-4

overview 30-4

configuration examples 30-18

default configuration 30-7

engine ID 30-7

groups 30-7, 30-10

host 30-7

ifIndex values 30-6

in-band management 1-7

in clusters 6-13

informs

and trap keyword 30-13

described 30-5

differences from traps 30-5

disabling 30-16

enabling 30-16

limiting access by TFTP servers 30-17

limiting system log messages to NMS 29-10

manager functions 1-6, 30-3

managing clusters with 6-16

notifications 30-5

overview 30-1, 30-5

security levels 30-3

setting CPU threshold notification 30-16

status, displaying 30-19

system contact and location 30-17

trap manager, configuring 30-14

traps

described 30-4, 30-5

differences from informs 30-5

disabling 30-16

enabling 30-13

enabling MAC address notification 5-16, 5-18, 5-19

overview 30-1, 30-5

types of 30-13

users 30-7, 30-10

versions supported 30-2

SNMP and Syslog Over IPv6 35-5

SNMPv1 30-2

SNMPv2C 30-3

SNMPv3 30-3

snooping, IGMP 21-2

software compatibility

See stacks, switch

software images

location in flash A-25

recovery procedures 38-2

scheduling reloads 3-22

tar file format, described A-25

See also downloading and uploading

source addresses

in IPv4 ACLs 31-10

source-and-destination-IP address based forwarding, EtherChannel 37-9

source-and-destination MAC address forwarding, EtherChannel 37-9

source-IP address based forwarding, EtherChannel 37-9

source-MAC address forwarding, EtherChannel 37-8

SPAN

and stack changes 27-10

configuration guidelines 27-11

default configuration 27-10

destination ports 27-7

displaying status 27-23

interaction with other features 27-9

monitored ports 27-6

- monitoring ports [27-7](#)
- overview [1-16, 27-1](#)
- ports, restrictions [23-12](#)
- received traffic [27-5](#)
- sessions
 - configuring ingress forwarding [27-15, 27-22](#)
 - creating [27-11](#)
 - defined [27-4](#)
 - limiting source traffic to specific VLANs [27-16](#)
 - removing destination (monitoring) ports [27-13](#)
 - specifying monitored ports [27-11](#)
 - with ingress traffic enabled [27-14](#)
- source ports [27-6](#)
- transmitted traffic [27-6](#)
- VLAN-based [27-7](#)
- spanning tree and native VLANs [13-14](#)
- Spanning Tree Protocol
 - See STP
- SPAN traffic [27-5](#)
- SRR
 - configuring
 - shaped weights on egress queues [33-77](#)
 - shared weights on egress queues [33-78](#)
 - shared weights on ingress queues [33-71](#)
 - described [33-14](#)
 - shaped mode [33-14](#)
 - shared mode [33-14](#)
 - support for [1-14, 1-15](#)
- SSH
 - configuring [9-43](#)
 - cryptographic software image [9-41](#)
 - described [1-7, 9-42](#)
 - encryption methods [9-42](#)
 - switch stack considerations [7-14](#)
 - user authentication methods, supported [9-42](#)
- SSL
 - configuration guidelines [9-48](#)
 - configuring a secure HTTP client [9-51](#)
 - configuring a secure HTTP server [9-50](#)
- cryptographic software image [9-46](#)
- described [9-46](#)
- monitoring [9-52](#)
- stack, switch
 - MAC address of [7-6, 7-17](#)
- stack changes, effects on
 - 802.1x port-based authentication [10-11](#)
 - ACL configuration [31-6](#)
 - CDP [25-2](#)
 - cross-stack EtherChannel [37-13](#)
 - EtherChannel [37-10](#)
 - IGMP snooping [21-6](#)
 - IP routing [34-3](#)
 - MAC address tables [5-15](#)
 - MSTP [17-8](#)
 - MVR [21-17](#)
 - port security [23-18](#)
 - SDM template selection [8-2](#)
 - SNMP [30-2](#)
 - SPAN and RSPAN [27-10](#)
 - STP [16-12](#)
 - switch clusters [6-13](#)
 - system message log [29-2](#)
 - VLANs [13-6](#)
 - VTP [14-8](#)
- stack master
 - bridge ID (MAC address) [7-6](#)
 - defined [7-1](#)
 - election [7-5](#)
 - IPv6 [35-6](#)
 - See also stacks, switch
- stack member
 - accessing CLI of specific member [7-21](#)
 - configuring
 - member number [7-19](#)
 - priority value [7-20](#)
 - defined [7-1](#)
 - displaying information of [7-22](#)
 - number [7-6](#)

- priority value [7-7](#)
- provisioning a new member [7-20](#)
- replacing [7-13](#)
- See also stacks, switch
- stack member number [12-16](#)
- stack protocol version [7-9](#)
- stacks, switch
 - accessing CLI of specific member [7-21](#)
 - assigning information
 - member number [7-19](#)
 - priority value [7-20](#)
 - provisioning a new member [7-20](#)
 - auto-advise [7-11](#)
 - auto-copy [7-10](#)
 - auto-extract [7-10](#)
 - auto-upgrade [7-10](#)
 - bridge ID [7-6](#)
 - CDP considerations [25-2](#)
 - compatibility, software [7-9](#)
 - configuration file [7-13](#)
 - configuration scenarios [7-15](#)
 - copying an image file from one member to another [A-38](#)
 - default configuration [7-17](#)
 - description of [7-1](#)
 - displaying information of [7-22](#)
 - enabling persistent MAC address timer [7-17](#)
 - in clusters [6-13](#)
 - incompatible software and image upgrades [7-13](#), [A-38](#)
 - IPv6 on [35-6](#)
 - MAC address considerations [5-15](#)
 - management connectivity [7-14](#)
 - managing [7-1](#)
 - membership [7-3](#)
 - merged [7-3](#)
 - MSTP instances supported [16-10](#)
 - offline configuration
 - described [7-7](#)
 - effects of adding a provisioned switch [7-7](#)
 - effects of removing a provisioned switch [7-9](#)
 - effects of replacing a provisioned switch [7-9](#)
 - provisioned configuration, defined [7-7](#)
 - provisioned switch, defined [7-7](#)
 - provisioning a new member [7-20](#)
 - partitioned [7-3](#), [38-8](#)
 - provisioned switch
 - adding [7-7](#)
 - removing [7-9](#)
 - replacing [7-9](#)
 - replacing a failed member [7-13](#)
 - software compatibility [7-9](#)
 - software image version [7-9](#)
 - stack protocol version [7-9](#)
 - STP
 - bridge ID [16-3](#)
 - instances supported [16-10](#)
 - root port selection [16-3](#)
 - stack root switch election [16-3](#)
 - system messages
 - hostnames in the display [29-1](#)
 - remotely monitoring [29-2](#)
 - system prompt consideration [5-8](#)
 - system-wide configuration considerations [7-13](#)
 - upgrading [A-38](#)
 - version-mismatch (VM) mode
 - automatic upgrades with auto-upgrade [7-10](#)
 - examples [7-11](#)
 - manual upgrades with auto-advise [7-11](#)
 - upgrades with auto-extract [7-10](#)
 - version-mismatch mode
 - described [7-10](#)
- See also stack master and stack member
- standby command switch
 - configuring
 - considerations [6-10](#)
 - defined [6-2](#)
 - priority [6-9](#)

- requirements [6-3](#)
- virtual IP address [6-10](#)
- See also cluster standby group and HSRP
- standby group, cluster
 - See cluster standby group and HSRP
- standby links [19-2](#)
- startup configuration
 - booting
 - manually [3-19](#)
 - specific image [3-19](#)
 - clearing [A-19](#)
 - configuration file
 - automatically downloading [3-18](#)
 - specifying the filename [3-18](#)
- static access ports
 - assigning to VLAN [13-9](#)
 - defined [12-3, 13-3](#)
- static addresses
 - See addresses
- static MAC addressing [1-10](#)
- static routes
 - configuring [34-5](#)
 - configuring for IPv6 [35-10](#)
- static VLAN membership [13-2](#)
- statistics
 - 802.1X [11-17](#)
 - 802.1x [10-65](#)
 - CDP [25-5](#)
 - interface [12-37](#)
 - LLDP [26-11](#)
 - LLDP-MED [26-11](#)
 - NMSP [26-11](#)
 - QoS ingress and egress [33-80](#)
 - RMON group Ethernet [28-5](#)
 - RMON group history [28-5](#)
 - SNMP input and output [30-19](#)
 - VTP [14-17](#)
- sticky learning [23-9](#)
- storm control
 - configuring [23-3](#)
 - described [23-1](#)
 - disabling [23-5](#)
 - displaying [23-20](#)
 - support for [1-4](#)
 - thresholds [23-2](#)
- STP
 - accelerating root port selection [18-4](#)
 - BackboneFast
 - described [18-7](#)
 - disabling [18-17](#)
 - enabling [18-17](#)
 - BPDU filtering
 - described [18-3](#)
 - disabling [18-15](#)
 - enabling [18-15](#)
 - BPDU guard
 - described [18-2](#)
 - disabling [18-14](#)
 - enabling [18-14](#)
 - BPDU message exchange [16-3](#)
 - configuration guidelines [16-13, 18-12](#)
 - configuring
 - forward-delay time [16-23](#)
 - hello time [16-22](#)
 - maximum aging time [16-23](#)
 - path cost [16-20](#)
 - port priority [16-18](#)
 - root switch [16-16](#)
 - secondary root switch [16-18](#)
 - spanning-tree mode [16-15](#)
 - switch priority [16-21](#)
 - transmit hold-count [16-24](#)
 - counters, clearing [16-24](#)
 - cross-stack UplinkFast
 - described [18-5](#)
 - enabling [18-16](#)
 - default configuration [16-13](#)
 - default optional feature configuration [18-12](#)

- designated port, defined 16-4
- designated switch, defined 16-4
- detecting indirect link failures 18-8
- disabling 16-16
- displaying status 16-24
- EtherChannel guard
 - described 18-10
 - disabling 18-18
 - enabling 18-17
- extended system ID
 - effects on root switch 16-16
 - effects on the secondary root switch 16-18
 - overview 16-4
 - unexpected behavior 16-16
- features supported 1-8
- IEEE 802.1D and bridge ID 16-4
- IEEE 802.1D and multicast addresses 16-9
- IEEE 802.1t and VLAN identifier 16-5
- inferior BPDU 16-3
- instances supported 16-10
- interface state, blocking to forwarding 18-2
- interface states
 - blocking 16-6
 - disabled 16-7
 - forwarding 16-6, 16-7
 - learning 16-7
 - listening 16-7
 - overview 16-5
- interoperability and compatibility among modes 16-11
- limitations with IEEE 802.1Q trunks 16-11
- load sharing
 - overview 13-19
 - using path costs 13-21
 - using port priorities 13-19
- loop guard
 - described 18-11
 - enabling 18-18
- modes supported 16-10
- multicast addresses, effect of 16-9
- optional features supported 1-9
- overview 16-2
- path costs 13-21, 13-22
- Port Fast
 - described 18-2
 - enabling 18-12
- port priorities 13-20
- preventing root switch selection 18-10
- protocols supported 16-10
- redundant connectivity 16-8
- root guard
 - described 18-10
 - enabling 18-18
- root port, defined 16-3
- root port selection on a switch stack 16-3
- root switch
 - configuring 16-16
 - effects of extended system ID 16-4, 16-16
 - election 16-3
 - unexpected behavior 16-16
- shutdown Port Fast-enabled port 18-2
- stack changes, effects of 16-12
- status, displaying 16-24
- superior BPDU 16-3
- timers, described 16-22
- UplinkFast
 - described 18-3
 - enabling 18-16
- stratum, NTP 5-3
- subnet mask 34-4
- success response, VMPS 13-23
- summer time 5-7
- SunNet Manager 1-6
- supported port-based authentication methods 10-7
- SVIs
 - and IP unicast routing 34-3
 - and router ACLs 31-4
 - connecting VLANs 12-10

- defined 12-3
 - switch 35-2
 - switch clustering technology 6-1
 - See also clusters, switch
 - switch console port 1-7
 - Switch Database Management
 - See SDM
 - Switched Port Analyzer
 - See SPAN
 - switched ports 12-2
 - switchport backup interface 19-4, 19-5
 - switchport block multicast command 23-8
 - switchport block unicast command 23-8
 - switchport protected command 23-7
 - switch priority
 - MSTP 17-22
 - STP 16-21
 - switch software features 1-1
 - switch virtual interface
 - See SVI
 - syslog
 - See system message logging
 - system capabilities TLV 26-2
 - system clock
 - configuring
 - daylight saving time 5-7
 - manually 5-5
 - summer time 5-7
 - time zones 5-6
 - displaying the time and date 5-5
 - overview 5-2
 - See also NTP
 - system description TLV 26-2
 - system message logging
 - default configuration 29-4
 - defining error message severity levels 29-9
 - disabling 29-4
 - displaying the configuration 29-14
 - enabling 29-5
 - facility keywords, described 29-14
 - level keywords, described 29-10
 - limiting messages 29-10
 - message format 29-2
 - overview 29-1
 - sequence numbers, enabling and disabling 29-8
 - setting the display destination device 29-5
 - stack changes, effects of 29-2
 - synchronizing log messages 29-6
 - syslog facility 1-16
 - time stamps, enabling and disabling 29-8
 - UNIX syslog servers
 - configuring the daemon 29-13
 - configuring the logging facility 29-13
 - facilities supported 29-14
 - system name
 - default configuration 5-9
 - default setting 5-9
 - manual configuration 5-9
 - See also DNS
 - system name TLV 26-2
 - system prompt, default setting 5-8, 5-9
 - system resources, optimizing 8-1
-
- ## T
- TACACS+
 - accounting, defined 9-11
 - authentication, defined 9-11
 - authorization, defined 9-11
 - configuring
 - accounting 9-17
 - authentication key 9-13
 - authorization 9-16
 - login authentication 9-14
 - default configuration 9-13
 - displaying the configuration 9-17
 - identifying the server 9-13
 - in clusters 6-15

- limiting the services to the user [9-16](#)
- operation of [9-12](#)
- overview [9-10](#)
- support for [1-12](#)
- tracking services accessed by user [9-17](#)
- tar files
 - creating [A-6](#)
 - displaying the contents of [A-6](#)
 - extracting [A-7](#)
 - image file format [A-25](#)
- TCAM
 - memory consistency check errors
 - example [38-26](#)
 - memory consistency check routines [1-5, 38-26](#)
 - memory consistency integrity [1-5, 38-26](#)
 - space
 - HFTM [38-26](#)
 - HQATM [38-26](#)
 - unassigned [38-26](#)
- TDR [1-16](#)
- Telnet
 - accessing management interfaces [2-10](#)
 - number of connections [1-7](#)
 - setting a password [9-6](#)
- temporary self-signed certificate [9-47](#)
- Terminal Access Controller Access Control System Plus
 - See TACACS+
- terminal lines, setting a password [9-6](#)
- ternary content addressable memory
 - See TCAM
- TFTP
 - configuration files
 - downloading [A-11](#)
 - preparing the server [A-10](#)
 - uploading [A-12](#)
 - configuration files in base directory [3-7](#)
 - configuring for autoconfiguration [3-7](#)
 - image files
 - deleting [A-28](#)
 - downloading [A-27](#)
 - preparing the server [A-26](#)
 - uploading [A-29](#)
 - limiting access by servers [30-17](#)
- TFTP server [1-6](#)
- threshold, traffic level [23-2](#)
- time
 - See NTP and system clock
- Time Domain Reflector
 - See TDR
- time-range command [31-15](#)
- time ranges in ACLs [31-15](#)
- time stamps in log messages [29-8](#)
- time zones [5-6](#)
- TLVs
 - defined [26-2](#)
 - LLDP [26-2](#)
 - LLDP-MED [26-2](#)
- Token Ring VLANs
 - support for [13-5](#)
 - VTP support [14-5](#)
- ToS [1-13](#)
- traceroute, Layer 2
 - and ARP [38-16](#)
 - and CDP [38-16](#)
 - broadcast traffic [38-15](#)
 - described [38-15](#)
 - IP addresses and subnets [38-16](#)
 - MAC addresses and VLANs [38-16](#)
 - multicast traffic [38-16](#)
 - multiple devices on a port [38-16](#)
 - unicast traffic [38-15](#)
 - usage guidelines [38-16](#)
- traceroute command [38-18](#)
 - See also IP traceroute
- traffic
 - blocking flooded [23-8](#)
 - fragmented [31-4](#)
 - unfragmented [31-4](#)

- traffic policing [1-14](#)
 - traffic suppression [23-2](#)
 - transmit hold-count
 - see STP
 - transparent mode, VTP [14-4](#)
 - trap-door mechanism [3-2](#)
 - traps
 - configuring MAC address notification [5-16, 5-18, 5-19](#)
 - configuring managers [30-13](#)
 - defined [30-4](#)
 - enabling [5-16, 5-18, 5-19, 30-13](#)
 - notification types [30-13](#)
 - overview [30-1, 30-5](#)
 - troubleshooting
 - connectivity problems [38-14, 38-15, 38-17](#)
 - CPU utilization [38-27](#)
 - detecting unidirectional links [24-1](#)
 - displaying crash information [38-23](#)
 - setting packet forwarding [38-21](#)
 - SFP security and identification [38-13](#)
 - show forward command [38-21](#)
 - with CiscoWorks [30-5](#)
 - with debug commands [38-19](#)
 - with ping [38-14](#)
 - with system message logging [29-1](#)
 - with traceroute [38-17](#)
 - trunk failover
 - See link-state tracking
 - trunking encapsulation [1-9](#)
 - trunk ports
 - configuring [13-16](#)
 - defined [12-3, 13-3](#)
 - trunks
 - allowed-VLAN list [13-17](#)
 - load sharing
 - setting STP path costs [13-21](#)
 - using STP port priorities [13-19, 13-20](#)
 - native VLAN for untagged traffic [13-18](#)
 - parallel [13-21](#)
 - pruning-eligible list [13-18](#)
 - to non-DTP device [13-13](#)
 - trusted boundary for QoS [33-45](#)
 - trusted port states
 - between QoS domains [33-47](#)
 - classification options [33-5](#)
 - ensuring port security for IP phones [33-45](#)
 - support for [1-14](#)
 - within a QoS domain [33-42](#)
 - trustpoints, CA [9-46](#)
 - twisted-pair Ethernet, detecting unidirectional links [24-1](#)
 - type of service
 - See ToS
-
- ## U
- UDLD
 - configuration guidelines [24-4](#)
 - default configuration [24-4](#)
 - disabling
 - globally [24-5](#)
 - on fiber-optic interfaces [24-5](#)
 - per interface [24-6](#)
 - echoing detection mechanism [24-3](#)
 - enabling
 - globally [24-5](#)
 - per interface [24-6](#)
 - link-detection mechanism [24-1](#)
 - neighbor database [24-2](#)
 - overview [24-1](#)
 - resetting an interface [24-6](#)
 - status, displaying [24-7](#)
 - support for [1-8](#)
 - unauthorized ports with IEEE 802.1x [10-10](#)
 - unicast MAC address filtering [1-6](#)
 - and adding static addresses [5-22](#)
 - and broadcast MAC addresses [5-21](#)
 - and CPU packets [5-21](#)

- and multicast addresses [5-21](#)
- and router MAC addresses [5-21](#)
- configuration guidelines [5-21](#)
- described [5-21](#)

unicast storm [23-1](#)

unicast storm control command [23-4](#)

unicast traffic, blocking [23-8](#)

UniDirectional Link Detection protocol

- See UDLD

UNIX syslog servers

- daemon configuration [29-13](#)
- facilities supported [29-14](#)
- message logging configuration [29-13](#)

unrecognized Type-Length-Value (TLV) support [14-5](#)

upgrading a Catalyst 2950 switch

- configuration compatibility issues [C-1](#)
- differences in configuration commands [C-1](#)
- feature behavior incompatibilities [C-5](#)
- incompatible command messages [C-1](#)
- recommendations [C-1](#)

upgrading software images

- See downloading

UplinkFast

- described [18-3](#)
- disabling [18-16](#)
- enabling [18-16](#)
- support for [1-8](#)

uploading

- configuration files
 - preparing [A-10](#), [A-13](#), [A-16](#)
 - reasons for [A-8](#)
 - using FTP [A-15](#)
 - using RCP [A-18](#)
 - using TFTP [A-12](#)
- image files
 - preparing [A-26](#), [A-30](#), [A-34](#)
 - reasons for [A-24](#)
 - using FTP [A-32](#)
 - using RCP [A-37](#)

- using TFTP [A-29](#)

USB mini-Type B console port [12-11](#)

USB Type A port [1-8](#)

user EXEC mode [2-2](#)

username-based authentication [9-6](#)

V

version-dependent transparent mode [14-5](#)

version-mismatch (VM) mode

- automatic upgrades with auto-upgrade [7-10](#)
- manual upgrades with auto-advise [7-11](#)
- upgrades with auto-extract [7-10](#)

version-mismatch mode

- described [7-10](#)

virtual IP address

- cluster standby group [6-10](#)
- command switch [6-10](#)

virtual switches and PAgP [37-6](#)

vlan.dat file [13-4](#)

VLAN 1, disabling on a trunk port [13-17](#)

VLAN 1 minimization [13-17](#)

vlan-assignment response, VMPS [13-23](#)

VLAN configuration

- at bootup [13-7](#)
- saving [13-7](#)

VLAN configuration mode [2-2](#)

VLAN database

- and startup configuration file [13-7](#)
- and VTP [14-1](#)
- VLAN configuration saved in [13-7](#)
- VLANs saved in [13-4](#)

VLAN filtering and SPAN [27-7](#)

vlan global configuration command [13-6](#)

VLAN ID, discovering [5-24](#)

VLAN load balancing on flex links [19-2](#)

- configuration guidelines [19-8](#)

VLAN management domain [14-2](#)

VLAN Management Policy Server

- See VMPS
- VLAN membership
 - confirming [13-26](#)
 - modes [13-3](#)
- VLAN Query Protocol
 - See VQP
- VLANs
 - adding [13-8](#)
 - adding to VLAN database [13-8](#)
 - aging dynamic addresses [16-9](#)
 - allowed on trunk [13-17](#)
 - and spanning-tree instances [13-3, 13-6, 13-11](#)
 - configuration guidelines, extended-range VLANs [13-11](#)
 - configuration guidelines, normal-range VLANs [13-6](#)
 - configuring [13-1](#)
 - configuring IDs 1006 to 4094 [13-11](#)
 - connecting through SVIs [12-10](#)
 - creating [13-8](#)
 - default configuration [13-7](#)
 - deleting [13-9](#)
 - described [12-2, 13-1](#)
 - displaying [13-13](#)
 - extended-range [13-1, 13-10](#)
 - features [1-9](#)
 - illustrated [13-2](#)
 - in the switch stack [13-6](#)
 - limiting source traffic with RSPAN [27-22](#)
 - limiting source traffic with SPAN [27-16](#)
 - modifying [13-8](#)
 - multicast [21-17](#)
 - native, configuring [13-18](#)
 - normal-range [13-1, 13-4](#)
 - number supported [1-9](#)
 - parameters [13-5](#)
 - port membership modes [13-3](#)
 - static-access ports [13-9](#)
 - STP and IEEE 802.1Q trunks [16-11](#)
 - supported [13-2](#)
 - Token Ring [13-5](#)
 - traffic between [13-2](#)
 - VTP modes [14-3](#)
- VLAN Trunking Protocol
 - See VTP
- VLAN trunks [13-13](#)
- VMPS
 - administering [13-27](#)
 - configuration example [13-28](#)
 - configuration guidelines [13-24](#)
 - default configuration [13-24](#)
 - description [13-23](#)
 - dynamic port membership
 - described [13-24](#)
 - reconfirming [13-26](#)
 - troubleshooting [13-28](#)
 - entering server address [13-25](#)
 - mapping MAC addresses to VLANs [13-23](#)
 - monitoring [13-27](#)
 - reconfirmation interval, changing [13-26](#)
 - reconfirming membership [13-26](#)
 - retry count, changing [13-27](#)
- voice aware 802.1x security
 - port-based authentication
 - configuring [10-39](#)
 - described [10-31, 10-39](#)
- voice-over-IP [15-1](#)
- voice VLAN
 - Cisco 7960 phone, port connections [15-1](#)
 - configuration guidelines [15-3](#)
 - configuring IP phones for data traffic
 - override CoS of incoming frame [15-6](#)
 - trust CoS priority of incoming frame [15-6](#)
 - configuring ports for voice traffic in
 - 802.1p priority tagged frames [15-5](#)
 - 802.1Q frames [15-5](#)
 - connecting to an IP phone [15-4](#)
 - default configuration [15-3](#)
 - described [15-1](#)

- displaying 15-7
 - IP phone data traffic, described 15-2
 - IP phone voice traffic, described 15-2
 - VQP 1-9, 13-23
 - VTP
 - adding a client to a domain 14-16
 - advertisements 13-15, 14-4
 - and extended-range VLANs 13-2, 14-2
 - and normal-range VLANs 13-2, 14-2
 - client mode, configuring 14-13
 - configuration
 - guidelines 14-9
 - requirements 14-11
 - saving 14-9
 - configuration requirements 14-11
 - configuration revision number
 - guideline 14-16
 - resetting 14-17
 - consistency checks 14-5
 - default configuration 14-9
 - described 14-1
 - domain names 14-9
 - domains 14-2
 - modes
 - client 14-3
 - off 14-4
 - server 14-3
 - transitions 14-3
 - transparent 14-4
 - monitoring 14-17
 - passwords 14-10
 - pruning
 - disabling 14-16
 - enabling 14-15
 - examples 14-7
 - overview 14-6
 - support for 1-9
 - pruning-eligible list, changing 13-18
 - server mode, configuring 14-11, 14-14
 - statistics 14-17
 - support for 1-9
 - Token Ring support 14-5
 - transparent mode, configuring 14-11
 - using 14-1
 - Version
 - enabling 14-14
 - version, guidelines 14-10
 - Version 1 14-5
 - Version 2
 - configuration guidelines 14-10
 - overview 14-5
 - Version 3
 - overview 14-5
-
- W
 - web authentication 10-17
 - configuring 11-16 to ??
 - described 1-10
 - web-based authentication
 - customizeable web pages 11-6
 - description 11-1
 - web-based authentication, interactions with other features 11-7
 - weighted tail drop
 - See WTD
 - wired location service
 - configuring 26-10
 - displaying 26-11
 - location TLV 26-3
 - understanding 26-4
 - wizards 1-2
 - WTD
 - described 33-13
 - setting thresholds
 - egress queue-sets 33-73
 - ingress queues 33-68
 - support for 1-14

X

Xmodem protocol [38-2](#)



Preface

Audience

This guide is for the networking professional managing the Catalyst 2960 and 2960-S switches, 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

This guide provides the information that you need to configure Cisco IOS software features on your switch. Catalyst 2960 and 2960-S switches run one of these images:

- The LAN base software image provides enterprise-class intelligent services such as access control lists (ACLs) and quality of service (QoS) features. On a Catalyst 2960-S switch, stacking is also supported.
- The LAN Lite image provides reduced functionality.

The Catalyst 2960-S ships with a universal image that includes cryptographic functionality. The software image on the switch is either the LAN base or LAN Lite image, depending on the switch model. To determine which image your switch is running:

- Switches running the LAN Lite image do not support the FlexStack module. They do not have a FlexStack module slot on the rear of the switch.
- On the front of the switch, the label in the top right corner ends in -S if the switch model runs the LAN Lite image.
- Enter the `show version` privileged EXEC command. The line that shows the product ID also ends in either -L (if running the LAN base image) or -S (if running the LAN Lite image). For example, WS-C2960S-48PD-L is running LAN base; WS-C2960S-24TS-S is running LAN Lite image.
- Enter the `show license` privileged EXEC command, and see which is the active image:

```
Switch# show license
Index 1 Feature: lanlite
      Period left: 0 minute 0 second
Index 2 Feature: lanbase
      Period left: Life time
      License Type: Permanent
      License State: Active, In Use
      License Priority: Medium
      License Count: Non-Counted
```

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, see the *Catalyst 2960 and 2960-S Switch Command Reference* for this release. For information about the standard Cisco IOS Release 12.4 commands, see the Cisco IOS documentation set available on Cisco.com.

This guide does not provide detailed information on the graphical user interfaces (GUIs) for the embedded device manager or for Cisco Network Assistant (hereafter referred to as *Network Assistant*) that you can use to manage the switch. However, the concepts in this guide are applicable to the GUI user. For information about the device manager, see the switch online help. For information about Network Assistant, see *Getting Started with Cisco Network Assistant*, available on Cisco.com.

This guide does not describe system messages you might encounter or how to install your switch. For more information, see the *Catalyst 2960 and 2960-S Switch System Message Guide* for this release and the *Catalyst 2960 and 2960-S Switch Hardware Installation Guide*.

For documentation updates, see the release notes for this release.

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.

Related Publications

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

http://www.cisco.com/en/US/products/ps6406/tsd_products_support_series_home.html

Catalyst 2960 Switch Getting Started Guide

Catalyst 2960-S Switch Getting Started Guide

Catalyst 2960 Hardware Installation Guide

Catalyst 2960-S Hardware Installation Guide



Note

Before installing, configuring, or upgrading the switch, see these documents:

- For initial configuration information, see the “Using Express Setup” section in the getting started guide or the “Configuring the Switch with the CLI-Based Setup Program” appendix in the hardware installation guide.
- For device manager requirements, see the “System Requirements” section in the release notes (not orderable but available on Cisco.com).
- For Network Assistant requirements, see the *Getting Started with Cisco Network Assistant* (not orderable but available on Cisco.com).
- For cluster requirements, see the *Release Notes for Cisco Network Assistant* (not orderable but available on Cisco.com).
- For upgrading information, see the “Downloading Software” section in the release notes.

See these documents for other information about the switch:

- *Release Notes for the Catalyst 3750, 3560, 2975, and 2960 Switches*
- *Catalyst 3750, 3560, 3550, 2975, 2975, 2970, and 2960 and 2960-S Switch System Message Guide*
- *Catalyst 2960 and 2960-S Switch Software Configuration Guide*
- *Catalyst 2960 and 2960-S Switch Command Reference*
- *Catalyst 2960 Switch Hardware Installation Guide*
- *Catalyst 2960-S Switch Hardware Installation Guide*
- *Catalyst 2960 Switch Getting Started Guide*
- *Catalyst 2960-S Switch Getting Started Guide*
- *Regulatory Compliance and Safety Information for the Catalyst 2960 and 2960-S Switches*
- *Auto Smartports Configuration Guide*
- *Cisco EnergyWise Configuration Guide*
- *Getting Started with Cisco Network Assistant*
- *Release Notes for Cisco Network Assistant*
- *Cisco RPS 300 Redundant Power System Hardware Installation Guide*
- *Cisco RPS 675 Redundant Power System Hardware Installation Guide*
- *Cisco Redundant Power System 2300 Hardware Installation Guide*

- For information about the Network Admission Control (NAC) features, see the *Network Admission Control Software Configuration Guide*
- Information about Cisco SFP, SFP+, and GBIC modules is available from this Cisco.com site: http://www.cisco.com/en/US/products/hw/modules/ps5455/prod_installation_guides_list.html
SFP compatibility matrix documents are available from this Cisco.com site: http://www.cisco.com/en/US/products/hw/modules/ps5455/products_device_support_tables_list.html

Obtaining Documentation, Obtaining Support, and Security Guidelines

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation:

<http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.



CHAPTER 1

Overview

This chapter provides these topics about the Catalyst 2960 and 2960-S switch software:

- [Features, page 1-1](#)
- [Default Settings After Initial Switch Configuration, page 1-16](#)
- [Network Configuration Examples, page 1-19](#)
- [Where to Go Next, page 1-24](#)

Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.

In this document, IP refers to IP Version 4 (IPv4) unless there is a specific reference to IP Version 6 (IPv6).

Features

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

- [Ease-of-Deployment and Ease-of-Use Features, page 1-2](#)
- [Performance Features, page 1-4](#)
- [Management Options, page 1-5](#)
- [Manageability Features, page 1-6](#)
- [Availability and Redundancy Features, page 1-8](#)
- [VLAN Features, page 1-9](#)
- [Security Features, page 1-10](#)
- [QoS and CoS Features, page 1-13](#)
- [Power over Ethernet Features, page 1-15](#)
- [Monitoring Features, page 1-16](#)

Ease-of-Deployment and Ease-of-Use Features

- Express Setup for quickly configuring a switch for the first time with basic IP information, contact information, switch and Telnet passwords, and Simple Network Management Protocol (SNMP) information through a browser-based program. For more information about Express Setup, see the getting started guide.
- User-defined and Cisco-default Smartports macros for creating custom switch configurations for simplified deployment across the network.
- An embedded device manager GUI for configuring and monitoring a single switch through a web browser. For information about launching the device manager, see the getting started guide. For more information about the device manager, see the switch online help.
- Cisco Network Assistant (hereafter referred to as *Network Assistant*) for
 - Managing communities, which are device groups like clusters, except that they can contain routers and access points and can be made more secure.
 - Simplifying and minimizing switch, switch stack, and switch cluster management from anywhere in your intranet.
 - Simplifying and minimizing switch and switch cluster management from anywhere in your intranet.
 - Accomplishing multiple configuration tasks from a single graphical interface without needing to remember command-line interface (CLI) commands to accomplish specific tasks.
 - Interactive guide mode that guides you in configuring complex features such as VLANs, ACLs, and quality of service (QoS).



Note If the switch is running the LAN Lite image, you can configure ACLs, but you cannot attach them to interfaces or VLANs.

- Configuration wizards that prompt you to provide only the minimum required information to configure complex features such as QoS priorities for traffic, priority levels for data applications, and security.
- Downloading an image to a switch.
- Applying actions to multiple ports and multiple switches at the same time, such as VLAN and QoS settings, inventory and statistic reports, link- and switch-level monitoring and troubleshooting, and multiple switch software upgrades.
- Viewing a topology of interconnected devices to identify existing switch clusters and eligible switches that can join a cluster and to identify link information between switches.
- Monitoring real-time status of a switch or multiple switches from the LEDs on the front-panel images. The system, redundant power system (RPS), and port LED colors on the images are similar to those used on the physical LEDs.



Note To use the RPS, the switch must be running the LAN Base image.



Note The Network Assistant must be downloaded from cisco.com/go/cna.

- Cisco FlexStack technology on Catalyst 2960-S switches running the LAN base image for
 - Connecting up to four switches through their FlexStack ports to operate as a single switch in the network.
 - Creating a bidirectional 20-Gb/s switching fabric across the switch stack, with all stack members having full access to the system bandwidth.
 - Using a single IP address and configuration file to manage the entire switch stack.
 - Automatic Cisco IOS version-check of new stack members with the option to automatically load images from the stack master or from a TFTP server.
 - Adding, removing, and replacing switches in the stack without disrupting the operation of the stack.
 - Provisioning a new member for a switch stack with the offline configuration feature. You can configure in advance the interface configuration for a specific stack member number and for a specific switch type of a new switch that is not part of the stack. The switch stack retains this information across stack reloads whether or not the provisioned switch is part of the stack.
 - Displaying stack-ring activity statistics (the number of frames sent by each stack member to the ring).
- Switch clustering technology for
 - Unified configuration, monitoring, authentication, and software upgrade of multiple, cluster-capable switches, regardless of their geographic proximity and interconnection media, including Ethernet, Fast Ethernet, Fast EtherChannel, small form-factor pluggable (SFP) modules, Gigabit Ethernet, and Gigabit EtherChannel connections. For a list of cluster-capable switches, see the release notes.
 - Automatic discovery of candidate switches and creation of clusters of up to 16 switches that can be managed through a single IP address.
 - Extended discovery of cluster candidates that are not directly connected to the command switch.
- Stack troubleshooting enhancements
- Auto Smartports
 - Cisco-default and user-defined macros for dynamic port configuration based on the device type detected on the port.
 - Enhancements to add support for global macros, last-resort macros, event trigger control, access points, EtherChannels, auto-QoS with Cisco Medianet, and IP phones.
 - Enhancements to add support for macro persistency, LLDP-based triggers, MAC address and OUI-based triggers, remote macros as well as for automatic configuration based on these two new device types: Cisco Digital Media Player (Cisco DMP) and Cisco IP Video Surveillance Camera (Cisco IPVSC).
 - Auto Smartports enhancement to enable auto-QoS on a CDP-capable Cisco digital media player.

For information, see the *Auto Smartports Configuration Guide*.

- Smart Install to allow a single point of management (director) in a network. You can use Smart Install to provide zero touch image and configuration upgrade of newly deployed switches and image and configuration downloads for any client switches. For more information, see the *Cisco Smart Install Configuration Guide*.

- Smart Install enhancements supporting client backup files, zero-touch replacement for clients with the same product-ID, automatic generation of the image list file, configurable file repository, hostname changes, transparent connection of the director to client, and USB storage for image and seed configuration.
- Smart Install enhancements in Cisco IOS Release 12.2(58)SE including the ability to manually change a client switch health state from denied to allowed or hold for on-demand upgrades, to remove selected clients from the director database, to allow simultaneous on-demand upgrade of multiple clients, and to provide more information about client devices, including device status, health status, and upgrade status.
- Call Home to provide e-mail-based and web-based notification of critical system events. Users with a service contract directly with Cisco Systems can register Call Home devices for the Cisco Smart Call Home service that generates automatic service requests with the Cisco TAC.

Performance Features

- Cisco EnergyWise manages the energy usage of endpoints connected to domain members. For more information, see the Cisco EnergyWise documentation on Cisco.com.
- EnergyWise Phase 2.5 enhancements that add support for a query to analyze and display domain information and for Wake on LAN (WoL) to remotely power on a WoL-capable PC.
- 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 and 10/100/1000 Mb/s interfaces and on 10/100/1000 BASE-TX SFP module interfaces that enables the interface to automatically detect the required cable connection type (straight-through or crossover) and to configure the connection appropriately.
- SFP+ support for 10Gigabit speeds (Catalyst 2960-S only)
- Support for up to 9000 bytes for frames that are bridged in hardware and up to 2000 bytes for frames that are bridged by software
- IEEE 802.3x flow control on all ports (the switch does not send pause frames).
- Up to 20 Gb/s of forwarding rates in a Catalyst 2960-S switch stack.
- EtherChannel for enhanced fault tolerance and for providing up to 8 Gb/s (Gigabit EtherChannel) or 800 Mb/s (Fast EtherChannel) full-duplex bandwidth among switches, routers, and servers.
- Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP) for automatic creation of EtherChannel links.
- Forwarding of Layer 2 packets at Gigabit line rate across the switches in the stack.
- Forwarding of Layer 2 packets at Gigabit line rate
- 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.
- Internet Group Management Protocol (IGMP) snooping for IGMP Versions 1, 2, and 3 for efficiently forwarding multimedia and multicast traffic
- IGMP report suppression for sending only one IGMP report per multicast router query to the multicast devices (supported only for IGMPv1 or IGMPv2 queries).

- IGMP snooping querier support to configure switch to generate periodic IGMP general query messages.
- IPv6 host support for basic IPv6 management
- Multicast Listener Discovery (MLD) snooping to enable efficient distribution of IP version 6 (IPv6) multicast data to clients and routers in a switched network



Note To use IPv6 features, the switch must be running the LAN Base image.

- 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.



Note To use MVR, the switch must be running the LAN Base image.

- IGMP filtering for controlling the set of multicast groups to which hosts on a switch port can belong.
- IGMP throttling for configuring the action when the maximum number of entries is in the IGMP forwarding table.
- IGMP leave timer for configuring the leave latency for the network.
- Switch Database Management (SDM) templates for allocating system resources to maximize support for user-selected features.
- Configurable small-frame arrival threshold to prevent storm control when small frames (64 bytes or less) arrive on an interface at a specified rate (the threshold).
- Flex Link Multicast Fast Convergence to reduce the multicast traffic convergence time after a Flex Link failure.



Note To use Flex Link Multicast Fast Convergence, the switch must be running the LAN Base image.

- RADIUS server load balancing to allow access and authentication requests to be distributed evenly across a server group.
- Support for QoS marking of CPU-generated traffic and queue CPU-generated traffic on the egress network ports.
- Memory consistency check routines to detect and correct invalid ternary content addressable memory (TCAM) table entries.

Management Options

- An embedded device manager—The device manager is a GUI that is integrated in the software image. You use it to configure and to monitor a single switch. For information about launching the device manager, see the getting started guide. For more information about the device manager, see the switch online help.
- Network Assistant—Network Assistant is a network management application that can be downloaded from Cisco.com. You use it to manage a single switch, a cluster of switches, or a community of devices. For more information about Network Assistant, see *Getting Started with Cisco Network Assistant*, available on Cisco.com.

- CLI—The Cisco IOS software supports desktop- and multilayer-switching features. You can access the CLI by connecting your management station directly to the switch console port, by connecting your PC directly to the Ethernet management port, or by using Telnet from a remote management station or PC. You can manage the switch stack by connecting to the console port or Ethernet management port of any stack member. For more information about the CLI, see [Chapter 1, “Using the Command-Line Interface.”](#)
- SNMP—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 1, “Configuring SNMP.”](#)
- Cisco IOS Configuration Engine (previously known to as the Cisco IOS CNS agent)—Configuration service automates the deployment and management of network devices and services. 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 CNS, see [Chapter 1, “Configuring Cisco IOS Configuration Engine.”](#)

Manageability Features

- CNS embedded agents for automating switch management, configuration storage, and delivery
- DHCP for automating configuration of switch information (such as IP address, default gateway, hostname, and Domain Name System [DNS] and TFTP server names)
- DHCP relay for forwarding User Datagram Protocol (UDP) broadcasts, including IP address requests, from DHCP clients
- DHCP server for automatic assignment of IP addresses and other DHCP options to IP hosts
- DHCP-based autoconfiguration and image update to download a specified configuration a new image to a large number of switches
- DHCP server port-based address allocation for the preassignment of an IP address to a switch port
- Directed unicast requests to a DNS server for identifying a switch through its IP address and its corresponding hostname 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
- Unicast MAC address filtering to drop packets with specific source or destination MAC addresses
- Configurable MAC address scaling that allows disabling MAC address learning on a VLAN to limit the size of the MAC address table
- Cisco Discovery Protocol (CDP) Versions 1 and 2 for network topology discovery and mapping between the switch and other Cisco devices on the network
- Link Layer Discovery Protocol (LLDP) and LLDP Media Endpoint Discovery (LLDP-MED) for interoperability with third-party IP phones
- LLDP media extensions (LLDP-MED) location TLV that provides location information from the switch to the endpoint device



Note To use LLDP-MED, the switch must be running the LAN Base image.

- Support for CDP and LLDP enhancements for exchanging location information with video end points for dynamic location-based content distribution from servers
- Network Time Protocol (NTP) version 4 for NTP time synchronization for both IPv4 and IPv6
- Cisco IOS File System (IFS) for providing a single interface to all file systems that the switch uses
- Configuration logging to log and to view changes to the switch configuration
- Unique device identifier to provide product identification information through a **show inventory** user EXEC command display
- In-band management access through the device manager over a Netscape Navigator or Microsoft Internet Explorer browser session
- 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 Secure Shell (SSH) connections for multiple CLI-based sessions over the network
- Support for SSH for IPv6.
- Support for IPv6 Host on the LAN Base and LAN Lite image (Catalyst 2960 and 2960-S).
- In-band management access through SNMP Versions 1, 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
- Out-of-band management access through the Ethernet management port to a PC (Catalyst 2960-only)
- Secure Copy Protocol (SCP) feature to provide a secure and authenticated method for copying switch configuration or switch image files (requires the cryptographic version of the software) for both IPv4 and IPv6
- Configuration replacement and rollback to replace the running configuration on a switch with any saved Cisco IOS configuration file
- The HTTP client in Cisco IOS supports can send requests to both IPv4 and IPv6 HTTP server, and the HTTP server in Cisco IOS can service HTTP requests from both IPv4 and IPv6 HTTP clients
- Simple Network and Management Protocol (SNMP) can be configured over IPv6 transport so that an IPv6 host can send SNMP queries and receive SNMP notifications from a device running IPv6
- IPv6 stateless autoconfiguration to manage link, subnet, and site addressing changes, such as management of host and mobile IP addresses
- Disabling MAC address learning on a VLAN
- DHCP server port-based address allocation for the preassignment of an IP address to a switch port.
- Wired location service sends location and attachment tracking information for connected devices to a Cisco Mobility Services Engine (MSE)



Note To use wired location, the switch must be running the LAN Base image.

- CPU utilization threshold trap monitors CPU utilization



Note To use CPU utilization, the switch must be running the LAN Base image.

- LLDP-MED network-policy profile time, length, value (TLV) for creating a profile for voice and voice-signalling by specifying the values for VLAN, class of service (CoS), differentiated services code point (DSCP), and tagging mode



Note Supported on all images in Cisco IOS Release 12.2(55)SE.

- Support for including a hostname in the option 12 field of DHCPDISCOVER packets. This provides identical configuration files to be sent by using the DHCP protocol
- DHCP Snooping enhancement to support the selection of a fixed string-based format for the circuit-id sub-option of the Option 82 DHCP field
- Increased support for LLDP-MED by allowing the switch to grant power to the power device (PD), based on the power policy TLV request
- USB mini-Type B console port in addition to the standard RJ-45 console port. Console input is active on only one port at a time. (Catalyst 2960-S only)
- USB Type A port for external Cisco USB flash memory devices (thumb drives or USB keys). You can use standard Cisco CLI commands to read, write, erase, copy, or boot from the flash memory. (Catalyst 2960-S only)

Availability and Redundancy Features

- Automatic stack master re-election for replacing stack masters that become unavailable (failover support)

The newly elected stack master begins accepting Layer 2 traffic in less than 1 second and Layer 3 traffic between 3 to 5 seconds. If traffic loss exceeds this time period, use the **stack-mac persistent timer** global configuration command to enable the persistent MAC address feature. When this feature is enabled, if the stack master changes, the stack MAC address does not change for a configured time value (1 to 60 minutes) or for an indefinite time period.

- Cross-stack EtherChannel for providing redundant links across the switch stack
- 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



Note Up to 64 spanning-tree instances are supported on switches running the LAN Lite image.

- Per-VLAN spanning-tree plus (PVST+) for load balancing across VLANs
- Rapid PVST+ for load balancing across VLANs and providing rapid convergence of spanning-tree instances

- UplinkFast, cross-stack UplinkFast, and BackboneFast for fast convergence after a spanning-tree topology change and for achieving load balancing between redundant uplinks, including Gigabit uplinks and cross-stack Gigabit uplinks
- 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 rapid per-VLAN Spanning-Tree plus (rapid-PVST+) based on the IEEE 802.1w Rapid Spanning Tree Protocol (RSTP) for rapid convergence of the spanning tree by immediately changing 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 change 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
- Flex Link Layer 2 interfaces to back up one another as an alternative to STP for basic link redundancy



Note To use Flex Links, the switch must be running the LAN Base image.

- Link-state tracking to mirror the state of the ports that carry upstream traffic from connected hosts and servers, and to allow the failover of the server traffic to an operational link on another Cisco Ethernet switch.



Note To use Link-state Tracking, the switch must be running the LAN Base image.

VLAN Features

- Support for up to 255 VLANs for assigning users to VLANs associated with appropriate network resources, traffic patterns, and bandwidth



Note Up to 64 VLANs are supported on switches running the LAN Lite image.

- Support for VLAN IDs in the 1 to 4094 range as allowed by the IEEE 802.1Q standard
- VLAN Query Protocol (VQP) for dynamic VLAN membership
- 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 (IEEE 802.1Q) 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. The switch CPU continues to send and receive control protocol frames.
- VLAN Flex Link Load Balancing to provide Layer 2 redundancy without requiring Spanning Tree Protocol (STP). A pair of interfaces configured as primary and backup links can load balance traffic based on VLAN.



Note To use VLAN Flex Link Load Balancing, the switch must be running the LAN Base image.

- Support for 802.1x authentication with restricted VLANs (also known as *authentication failed VLANs*)
- Support for VTP version 3 that includes support for configuring extended range VLANs (VLANs 1006 to 4094) in any VTP mode, enhanced authentication (hidden or secret passwords), propagation of other databases in addition to VTP, VTP primary and secondary servers, and the option to turn VTP on or off by port

Security Features

- Web authentication to allow a supplicant (client) that does not support IEEE 802.1x functionality to be authenticated using a web browser



Note To use Web Authentication, the switch must be running the LAN Base image.

- Local web authentication banner so that a custom banner or an image file can be displayed at a web authentication login screen
- IEEE 802.1x Authentication with ACLs and the RADIUS Filter-Id Attribute



Note To use this feature, the switch must be running the LAN Base image.

- Password-protected access (read-only and read-write access) to management interfaces (device manager, Network Assistant, and the CLI) 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
- VLAN aware port security option to shut down the VLAN on the port when a violation occurs, instead of shutting down the entire port.

- Port security aging to set the aging time for secure addresses on a port
- Protocol storm protection to control the rate of incoming protocol traffic to a switch by dropping packets that exceed a specified ingress rate.
- 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 inbound security policies on Layer 2 interfaces (port ACLs)
- Extended MAC access control lists for defining security policies in the inbound direction on Layer 2 interfaces
- Source and destination MAC-based ACLs for filtering non-IP traffic
- DHCP snooping to filter untrusted DHCP messages between untrusted hosts and DHCP servers
- IP source guard to restrict traffic on nonrouted interfaces by filtering traffic based on the DHCP snooping database and IP source bindings
- Dynamic ARP inspection to prevent malicious attacks on the switch by not relaying invalid ARP requests and responses to other ports in the same VLAN
- IEEE 802.1x port-based authentication to prevent unauthorized devices (clients) from gaining access to the network. These features are supported:
 - Multidomain authentication (MDA) to allow both a data device and a voice device, such as an IP phone (Cisco or non-Cisco), to independently authenticate on the same IEEE 802.1x-enabled switch port



Note To use MDA, the switch must be running the LAN Base image.

- Dynamic voice virtual LAN (VLAN) for MDA to allow a dynamic voice VLAN on an MDA-enabled port
- VLAN assignment for restricting 802.1x-authenticated users to a specified VLAN
- Support for VLAN assignment on a port configured for multi-auth mode. The RADIUS server assigns a VLAN to the first host to authenticate on the port, and subsequent hosts use the same VLAN. Voice VLAN assignment is supported for one IP phone.



Note To use this feature, the switch must be running the LAN Base image.

- Port security for controlling access to 802.1x ports
- Voice VLAN to permit a Cisco IP Phone to access the voice VLAN regardless of the authorized or unauthorized state of the port
- IP phone detection enhancement to detect and recognize a Cisco IP phone.
- Guest VLAN to provide limited services to non-802.1x-compliant users
- Restricted VLAN to provide limited services to users who are 802.1x compliant, but do not have the credentials to authenticate via the standard 802.1x processes



Note To use authentication with restricted VLANs, the switch must be running the LAN Base image.

- 802.1x accounting to track network usage

- 802.1x with wake-on-LAN to allow dormant PCs to be powered on based on the receipt of a specific Ethernet frame
- 802.1x readiness check to determine the readiness of connected end hosts before configuring IEEE 802.1x on the switch



Note To use 802.1x readiness check, the switch must be running the LAN Base image.

- Voice aware 802.1x security to apply traffic violation actions only on the VLAN on which a security violation occurs.



Note To use voice aware 802.1x authentication, the switch must be running the LAN Base image.

- MAC authentication bypass to authorize clients based on the client MAC address.



Note To use MAC authentication bypass, the switch must be running the LAN Base image.

- Network Admission Control (NAC) Layer 2 802.1x validation of the antivirus condition or *posture* of endpoint systems or clients before granting the devices network access.

For information about configuring NAC Layer 2 802.1x validation, see the [“Configuring NAC Layer 2 802.1x Validation” section on page 1-58](#).



Note To use NAC, the switch must be running the LAN Base image.

- Network Edge Access Topology (NEAT) with 802.1X switch supplicant, host authorization with CISP, and auto enablement to authenticate a switch outside a wiring closet as a supplicant to another switch.
- IEEE 802.1x with open access to allow a host to access the network before being authenticated.
- IEEE 802.1x authentication with downloadable ACLs and redirect URLs to allow per-user ACL downloads from a Cisco Secure ACS server to an authenticated switch.
- Support for dynamic creation or attachment of an auth-default ACL on a port that has no configured static ACLs.



Note To use this feature, the switch must be running the LAN Base image.

- Flexible-authentication sequencing to configure the order of the authentication methods that a port tries when authenticating a new host.
- Multiple-user authentication to allow more than one host to authenticate on an 802.1x-enabled port.
- TACACS+, a proprietary feature for managing network security through a TACACS server for both IPv4 and IPv6
- RADIUS for verifying the identity of, granting access to, and tracking the actions of remote users through authentication, authorization, and accounting (AAA) services for both IPv4 and IPv6
- Enhancements to RADIUS, TACACS+, and SSH to function over IPv6.

- Secure Socket Layer (SSL) Version 3.0 support for the HTTP 1.1 server authentication, encryption, and message integrity and HTTP client authentication to allow secure HTTP communications (requires the cryptographic version of the software)
- IEEE 802.1x Authentication with ACLs and the RADIUS Filter-Id Attribute
- Support for IP source guard on static hosts.
- RADIUS Change of Authorization (CoA) to change the attributes of a certain session after it is authenticated. When there is a change in policy for a user or user group in AAA, administrators can send the RADIUS CoA packets from the AAA server, such as Cisco Secure ACS to reinitialize authentication, and apply to the new policies.
- IEEE 802.1x User Distribution to allow deployments with multiple VLANs (for a group of users) to improve scalability of the network by load balancing users across different VLANs. Authorized users are assigned to the least populated VLAN in the group, assigned by RADIUS server.
- Support for critical VLAN with multiple-host authentication so that when a port is configured for multi-auth, and an AAA server becomes unreachable, the port is placed in a critical VLAN in order to still permit access to critical resources.
- Customizable web authentication enhancement to allow the creation of user-defined *login*, *success*, *failure* and *expire* web pages for local web authentication.
- Support for Network Edge Access Topology (NEAT) to change the port host mode and to apply a standard port configuration on the authenticator switch port.
- VLAN-ID based MAC authentication to use the combined VLAN and MAC address information for user authentication to prevent network access from unauthorized VLANs.
- MAC move to allow hosts (including the hosts connected behind an IP phone) to move across ports within the same switch without any restrictions to enable mobility. With MAC move, the switch treats the reappearance of the same MAC address on another port in the same way as a completely new MAC address.
- Support for 3DES and AES with version 3 of the Simple Network Management Protocol (SNMPv3). This release adds support for the 168-bit Triple Data Encryption Standard (3DES) and the 128-bit, 192-bit, and 256-bit Advanced Encryption Standard (AES) encryption algorithms to SNMPv3.

QoS and CoS Features

- Automatic QoS (auto-QoS) to simplify the deployment of existing QoS features by classifying traffic and configuring egress queues



Note To use auto-QoS, the switch must be running the LAN Base image.

- Classification
 - IP type-of-service/Differentiated Services Code Point (IP ToS/DSCP) and IEEE 802.1p CoS marking priorities on a per-port basis for protecting the performance of mission-critical applications



Note To use DSCP, the switch must be running the LAN Base image.

- IP ToS/DSCP and IEEE 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



Note To use flow-based packet classification, the switch must be running the LAN Base image.

- 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



Note To use trusted boundary, the switch must be running the LAN Base image.

- Policing



Note To use policy maps, the switch must be running the LAN Base image

- Traffic-policing policies on the switch port for managing how much of the port bandwidth should be allocated to a specific traffic flow
- If you configure multiple class maps for a hierarchical policy map, each class map can be associated with its own port-level (second-level) policy map. Each second-level policy map can have a different policer.
- Aggregate policing for policing traffic flows in aggregate to restrict specific applications or traffic flows to metered, predefined rates
- Out-of-Profile
 - Out-of-profile markdown for packets that exceed bandwidth utilization limits
- 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



Note To use WTD, the switch must be running the LAN Base image.

- Shaped round robin (SRR) as the scheduling service for specifying the rate at which packets are sent to the internal ring (sharing is the only supported mode on ingress queues)



Note To use ingress queueing, the switch must be running the LAN Base image.



Note Ingress queueing is not supported on Catalyst 2960-S switches.

- 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 specifying 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.



Note To use egress queuing, the switch must be running the LAN Base image.

- Auto-QoS enhancements that add automatic configuration classification of traffic flow from video devices, such as the Cisco Telepresence System and Cisco Surveillance Camera.



Note To use Auto-QoS enhancements, the switch must be running the LAN Base image.

Layer 3 Features

- When you configure the **lanbase-routing** SDM template, the switch supports static routing and router ACLs on SVIs (supported only on switches running the LAN base image).
- IPv6 default router preference (DRP) for improving the ability of a host to select an appropriate router (requires the LAN Base image)

Power over Ethernet Features

- Ability to provide power to connected Cisco pre-standard and IEEE 802.3af-compliant powered devices from Power over Ethernet (PoE)-capable ports if the switch detects that there is no power on the circuit.
- Support for IEEE 802.3at, (PoE+) that increases the available power that can be drawn by powered devices from 15.4 W per port to 30 W per port (Catalyst 2960-S only)
- Support for CDP with power consumption. The powered device notifies the switch of the amount of power it is consuming.
- Support for Cisco intelligent power management. The powered device and the switch negotiate through power-negotiation CDP messages for an agreed power-consumption level. The negotiation allows a high-power Cisco powered device to operate at its highest power mode.
- Automatic detection and power budgeting; the switch maintains a power budget, monitors and tracks requests for power, and grants power only when it is available.
- Ability to monitor the real-time power consumption. On a per-PoE port basis, the switch senses the total power consumption, polices the power usage, and reports the power usage.

Monitoring Features

- Switch LEDs that provide port- and switch-level status
- Switch LEDs that provide port-, switch-, and stack-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 port or VLAN
- 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
- Time Domain Reflector (TDR) to diagnose and resolve cabling problems on 10/100 and 10/100/1000 copper Ethernet ports
- SFP module diagnostic management interface to monitor physical or operational status of an SFP module
- Generic online diagnostics to test hardware functionality of the supervisor engine, modules, and switch while the switch is connected to a live network (Catalyst 2960-S only).
- On-board failure logging (OBFL) to collect information about the switch and the power supplies connected to it (Catalyst 2960-S only)
- IP Service Level Agreements (IP SLAs) responder support that allows the switch to be a target device for IP SLAs active traffic monitoring



Note To use IP SLAs, the switch must be running the LAN Base image.

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- and stack-wide settings.



Note

For information about assigning an IP address by using the browser-based Express Setup program, see the getting started guide. For information about assigning an IP address by using the CLI-based setup program, see the hardware installation guide.

If you do not configure the switch at all, the switch operates with these default settings:

- Default switch IP address, subnet mask, and default gateway is 0.0.0.0. For more information, see [Chapter 1, “Assigning the Switch IP Address and Default Gateway,”](#) and [Chapter 1, “Configuring DHCP and IP Source Guard Features.”](#)
- Default domain name is not configured. For more information, see [Chapter 1, “Assigning the Switch IP Address and Default Gateway.”](#)
- DHCP client is enabled, the DHCP server is enabled (only if the device acting as a DHCP server is configured and is enabled), and the DHCP relay agent is enabled (only if the device is acting as a DHCP relay agent is configured and is enabled). For more information, see [Chapter 1, “Assigning the Switch IP Address and Default Gateway,”](#) and [Chapter 1, “Configuring DHCP and IP Source Guard Features.”](#)
- Switch cluster is disabled. For more information about switch clusters, see [Chapter 1, “Clustering Switches,”](#) and the *Getting Started with Cisco Network Assistant*, available on Cisco.com.
- No passwords are defined. For more information, see [Chapter 1, “Administering the Switch.”](#)
- System name and prompt is *Switch*. For more information, see [Chapter 1, “Administering the Switch.”](#)
- NTP is enabled. For more information, see [Chapter 1, “Administering the Switch.”](#)
- DNS is enabled. For more information, see [Chapter 1, “Administering the Switch.”](#)
- TACACS+ is disabled. For more information, see [Chapter 1, “Configuring Switch-Based Authentication.”](#)
- RADIUS is disabled. For more information, see [Chapter 1, “Configuring Switch-Based Authentication.”](#)
- The standard HTTP server and Secure Socket Layer (SSL) HTTPS server are both enabled. For more information, see [Chapter 1, “Configuring Switch-Based Authentication.”](#)
- IEEE 802.1x is disabled. For more information, see [Chapter 1, “Configuring IEEE 802.1x Port-Based Authentication.”](#)
- Port parameters
 - Interface speed and duplex mode is autonegotiate. For more information, see
 - Auto-MDIX is enabled. For more information, see
 - Flow control is off. For more information, see
 - PoE is autonegotiate. For more information, see
- VLANs
 - Default VLAN is VLAN 1. For more information, see [Chapter 1, “Configuring VLANs.”](#)
 - VLAN trunking setting is dynamic auto (DTP). For more information, see [Chapter 1, “Configuring VLANs.”](#)
 - Trunk encapsulation is negotiate. For more information, see [Chapter 1, “Configuring VLANs.”](#)
 - VTP mode is server. For more information, see [Chapter 1, “Configuring VTP.”](#)
 - VTP version is Version 1. For more information, see [Chapter 1, “Configuring VTP.”](#)
 - Voice VLAN is disabled. For more information, see [Chapter 1, “Configuring Voice VLAN.”](#)
- STP, PVST+ is enabled on VLAN 1. For more information, see [Chapter 1, “Configuring STP.”](#)
- MSTP is disabled. For more information, see [Chapter 1, “Configuring MSTP.”](#)

- Optional spanning-tree features are disabled. For more information, see [Chapter 1, “Configuring Optional Spanning-Tree Features.”](#)
- Flex Links are not configured. For more information, see [Chapter 1, “Configuring Flex Links and the MAC Address-Table Move Update Feature.”](#)



Note To use Flex Links, the switch must be running the LAN Base image.

- DHCP snooping is disabled. The DHCP snooping information option is enabled. For more information, see [Chapter 1, “Configuring DHCP and IP Source Guard Features.”](#)
- IP source guard is disabled. For more information, see [Chapter 1, “Configuring DHCP and IP Source Guard Features.”](#)
- DHCP server port-based address allocation is disabled. For more information, see [Chapter 1, “Configuring DHCP and IP Source Guard Features.”](#)
- Dynamic ARP inspection is disabled on all VLANs. For more information, see [Chapter 1, “Configuring Dynamic ARP Inspection.”](#)
- IGMP snooping is enabled. No IGMP filters are applied. For more information, see [Chapter 1, “Configuring IGMP Snooping and MVR.”](#)
- IGMP throttling setting is deny. For more information, see [Chapter 1, “Configuring IGMP Snooping and MVR.”](#)
- The IGMP snooping querier feature is disabled. For more information, see [Chapter 1, “Configuring IGMP Snooping and MVR.”](#)
- MVR is disabled. For more information, see [Chapter 1, “Configuring IGMP Snooping and MVR.”](#)



Note To use MVR, the switch must be running the LAN Base image.

- Port-based traffic
 - Broadcast, multicast, and unicast storm control is disabled. For more information, see [Chapter 1, “Configuring Port-Based Traffic Control.”](#)
 - No protected ports are defined. For more information, see [Chapter 1, “Configuring Port-Based Traffic Control.”](#)
 - Unicast and multicast traffic flooding is not blocked. For more information, see [Chapter 1, “Configuring Port-Based Traffic Control.”](#)
 - No secure ports are configured. For more information, see [Chapter 1, “Configuring Port-Based Traffic Control.”](#)
- CDP is enabled. For more information, see [Chapter 1, “Configuring CDP.”](#)
- UDLD is disabled. For more information, see [Chapter 1, “Configuring UDLD.”](#)
- SPAN and RSPAN are disabled. For more information, see [Chapter 1, “Configuring SPAN and RSPAN.”](#)



Note To use RSPAN, the switch must be running the LAN Base image.

- RMON is disabled. For more information, see [Chapter 1, “Configuring RMON.”](#)

- Syslog messages are enabled and appear on the console. For more information, see [Chapter 1, “Configuring System Message Logging.”](#)
- SNMP is enabled (Version 1). For more information, see [Chapter 1, “Configuring SNMP.”](#)
- No ACLs are configured. For more information, see [Chapter 1, “Configuring Network Security with ACLs.”](#)
- QoS is disabled. For more information, see [Chapter 1, “Configuring QoS.”](#)
- No EtherChannels are configured. For more information, see [Chapter 1, “Configuring EtherChannels and Link-State Tracking.”](#)

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.

- [“Design Concepts for Using the Switch” section on page 1-19](#)
- [“Small to Medium-Sized Network Using Catalyst 2960 and 2960-S Switches” section on page 1-22](#)
- [“Long-Distance, High-Bandwidth Transport Configuration” section on page 1-23](#)

Design Concepts for Using the Switch

As your network users compete for network bandwidth, it takes longer to send and receive data. When you configure your network, consider the bandwidth required by your network users and the relative priority of the network applications that they use.

[Table 1-1](#) describes what can cause network performance to degrade and how you can configure your network to increase the bandwidth available to your network users.

Table 1-1 *Increasing Network Performance*

Network Demands	Suggested Design Methods
Too many users on a single network segment and a growing number of users accessing the Internet	<ul style="list-style-type: none"> • Create smaller network segments so that fewer users share the bandwidth, and use VLANs and IP subnets to place the network resources in the same logical network as the users who access those resources most. • Use full-duplex operation between the switch and its connected workstations.
<ul style="list-style-type: none"> • Increased power of new PCs, workstations, and servers • High bandwidth demand from networked applications (such as e-mail with large attached files) and from bandwidth-intensive applications (such as multimedia) 	<ul style="list-style-type: none"> • Connect global resources—such as servers and routers to which the network users require equal access—directly to the high-speed switch ports so that they have their own high-speed segment. • Use the EtherChannel feature between the switch and its connected servers and routers.

Bandwidth alone is not the only consideration when designing your network. As your network traffic profiles evolve, consider providing network services that can support applications for voice and data integration, multimedia integration, application prioritization, and security. Table 1-2 describes some network demands and how you can meet them.

Table 1-2 Providing Network Services

Network Demands	Suggested Design Methods
Efficient bandwidth usage for multimedia applications and guaranteed bandwidth for critical applications	<ul style="list-style-type: none"> Use IGMP snooping to efficiently forward multimedia and multicast traffic. Use other QoS mechanisms such as packet classification, marking, scheduling, and congestion avoidance to classify traffic with the appropriate priority level, thereby providing maximum flexibility and support for mission-critical, unicast, and multicast and multimedia applications. Use MVR to continuously send multicast streams in a multicast VLAN but to isolate the streams from subscriber VLANs for bandwidth and security reasons.
High demand on network redundancy and availability to provide <i>always on</i> mission-critical applications	<ul style="list-style-type: none"> Use switch stacks, where all stack members are eligible stack masters in case of stack-master failure. All stack members have synchronized copies of the saved and running configuration files of the switch stack. <p>Note Stacking is supported only on Catalyst 2960-S switches running the LAN base image.</p> <ul style="list-style-type: none"> Use cross-stack EtherChannels for providing redundant links across the switch stack. Use VLAN trunks, cross-stack UplinkFast, and BackboneFast for traffic-load balancing on the uplink ports so that the uplink port with a lower relative port cost is selected to carry the VLAN traffic.
High demand on network redundancy and availability to provide <i>always on</i> mission-critical applications	<ul style="list-style-type: none"> Use VLAN trunks and BackboneFast for traffic-load balancing on the uplink ports so that the uplink port with a lower relative port cost is selected to carry the VLAN traffic.
An evolving demand for IP telephony	<ul style="list-style-type: none"> Use QoS to prioritize applications such as IP telephony during congestion and to help control both delay and jitter within the network. Use switches that support at least two queues per port to prioritize voice and data traffic as either high- or low-priority, based on IEEE 802.1p/Q. The switch supports at least four queues per port. Use voice VLAN IDs (VVIDs) to provide separate VLANs for voice traffic.
A growing demand for using existing infrastructure to transport data and voice from a home or office to the Internet or an intranet at higher speeds	<p>Use the Catalyst Long-Reach Ethernet (LRE) switches to provide up to 15 Mb of IP connectivity over existing infrastructure, such as existing telephone lines.</p> <p>Note To use LRE, the switch must be running the LAN Base image.</p> <p>Note LRE is the technology used in the Catalyst 2900 LRE XL and Catalyst 2950 LRE switches. See the documentation sets specific to these switches for LRE information.</p>

You can use the switches and switch stacks to create the following:

- Catalyst 2960-S switches. To preserve switch connectivity if one switch in the stack fails, connect the switches as recommended in the hardware installation guide, and enable either cross-stack Etherchannel or cross-stack UplinkFast.


You can have redundant uplink connections, using SFP modules in the switch stack to a Gigabit backbone switch, such as a Catalyst 4500 or Catalyst 3750-12S Gigabit switch. You can also create backup paths by using Fast Ethernet, Gigabit, or EtherChannel links. If one of the redundant connections fails, the other can serve as a backup path. If the Gigabit switch is cluster-capable, you can configure it and the switch stack as a switch cluster to manage them through a single IP address. The Gigabit switch can be connected to a Gigabit server through a 1000BASE-T connection.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

Figure 1-1 Cost-Effective Wiring Closet



- Server aggregation (Figure 1-2)—You can use the switches to interconnect groups of servers, centralizing physical security and administration of your network. For high-speed IP forwarding at the distribution layer, connect the switches in the access layer to multilayer switches with routing capability. The Gigabit interconnections minimize latency in the data flow.


QoS and policing on the switches provide preferential treatment for certain data streams. They segment traffic streams into different paths for processing. Security features on the switch ensure rapid handling of packets.

Fault tolerance from the server racks to the core is achieved through dual homing of servers connected to switches, which have redundant Gigabit EtherChannels.

Using dual SFP module uplinks from the switches provides redundant uplinks to the network core. Using SFP modules provides flexibility in media and distance options through fiber-optic connections.

The various lengths of stack cable available, ranging from 0.5 meter to 3 meters provide extended connections to the switch stacks across multiple server racks, for multiple stack aggregation.

Figure 1-2 Server Aggregation



Small to Medium-Sized Network Using Catalyst 2960 and 2960-S Switches

Figure 1-3 shows a configuration for a network of up to 500 employees. This network uses The switches are using EtherChannel for load sharing.

The switches are connected to workstations and local servers. The server farm includes a call-processing server running Cisco CallManager software. Cisco CallManager controls call processing, routing, and Cisco IP Phone features and configuration. The switches are interconnected through Gigabit interfaces.

This network uses VLANs to logically segment the network into well-defined broadcast groups and for security management. Data and multimedia traffic are configured on the same VLAN. Voice traffic from the Cisco IP Phones are configured on separate VVIDs. If data, multimedia, and voice traffic are assigned to the same VLAN, only one VLAN can be configured per wiring closet.


When an end station in one VLAN needs to communicate with an end station in another VLAN, a router or Layer 3 switch routes the traffic to the destination VLAN. In this network, the routers are providing inter-VLAN routing. VLAN access control lists (VLAN maps) on the switch provide intra-VLAN security and prevent unauthorized users from accessing critical areas of the network.

In addition to inter-VLAN routing, the routers provide QoS mechanisms such as DSCP priorities to prioritize the different types of network traffic and to deliver high-priority traffic. If congestion occurs, QoS drops low-priority traffic to allow delivery of high-priority traffic.

Cisco CallManager controls call processing, routing, and Cisco IP Phone features and configuration. Users with workstations running Cisco SoftPhone software can place, receive, and control calls from their PCs. Using Cisco IP Phones, Cisco CallManager software, and Cisco SoftPhone software integrates telephony and IP networks, and the IP network supports both voice and data.

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

Figure 1-3 Collapsed Backbone Configuration



Long-Distance, High-Bandwidth Transport Configuration



Note

To use CWDM SFPs, the switch must be running the LAN Base image.

Figure 1-4 shows a configuration for sending 8 Gigabits of data over a single fiber-optic cable. The Catalyst 2960 or 2960-S switches have coarse wavelength-division multiplexing (CWDM) fiber-optic SFP modules installed. Depending on the CWDM SFP module, data is sent at wavelengths from 1470 to 1610 nm. The higher the wavelength, the farther the transmission can travel. A common wavelength used for long-distance transmissions is 1550 nm.

The CWDM SFP modules connect to CWDM optical add/drop multiplexer (OADM) modules over distances of up to 393,701 feet (74.5 miles or 120 km). The CWDM OADM modules combine (or *multiplex*) the different CWDM wavelengths, allowing them to travel simultaneously on the same fiber-optic cable. The CWDM OADM modules on the receiving end separate (or *demultiplex*) the different wavelengths.

For more information about the CWDM SFP modules and CWDM OADM modules, see the *Cisco CWDM GBIC and CWDM SFP Installation Note*.

Figure 1-4 Long-Distance, High-Bandwidth Transport Configuration

Where to Go Next

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

- [Chapter 1, “Using the Command-Line Interface”](#)
- [Chapter 1, “Assigning the Switch IP Address and Default Gateway”](#)

To locate and download MIBs for a specific Cisco product and release, use the Cisco MIB Locator:
<http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>.



CHAPTER 1

Using the Command-Line Interface

This chapter describes the Cisco IOS command-line interface (CLI) and how to use it to configure your Catalyst 2960 or 2960-S switch. Unless otherwise noted, the term switch refers to a standalone switch and to a switch stack.

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

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 1-1 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 hostname *Switch*.

Table 1-1 Command Mode Summary

Mode	Access Method	Prompt	Exit Method	About This Mode
User EXEC	Begin a session with your switch.	Switch>	Enter logout or quit .	Use this mode to <ul style="list-style-type: none"> • Change terminal settings. • Perform basic tests. • Display system information.
Privileged EXEC	While in user EXEC mode, enter the enable command.	Switch#	Enter disable to exit.	Use this mode to verify commands that you have entered. Use a password to protect access to this mode.
Global configuration	While in privileged EXEC mode, enter the configure command.	Switch(config)#	To exit to privileged EXEC mode, enter exit or end , or press Ctrl-Z .	Use this mode to configure parameters that apply to the entire switch.
VLAN configuration	While in global configuration mode, enter the vlan <i>vlan-id</i> 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.
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 1-15. To configure multiple interfaces with the same parameters, see the “Configuring a Range of Interfaces” section on page 1-17.
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.

For more detailed information on the command modes, see the command reference guide for this release.

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 1-2](#).

Table 1-2 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: <pre>Switch# di? dir disable disconnect</pre>
<i>abbreviated-command-entry<Tab></i>	Complete a partial command name. For example: <pre>Switch# sh conf<tab> Switch# show configuration</pre>
?	List all commands available for a particular command mode. For example: <pre>Switch> ?</pre>
<i>command ?</i>	List the associated keywords for a command. For example: <pre>Switch> show ?</pre>
<i>command keyword ?</i>	List the associated arguments for a keyword. For example: <pre>Switch(config)# cdp holdtime ? <10-255> Length of time (in sec) that receiver must keep this packet</pre>

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 1-3 lists some error messages that you might encounter while using the CLI to configure your switch.

Table 1-3 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 appear.
% 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 appear.
% 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 appear.

Using Configuration Logging

You can log and view changes to the switch configuration. You can use the Configuration Change Logging and Notification feature to track changes on a per-session and per-user basis. The logger tracks each configuration command that is applied, the user who entered the command, the time that the command was entered, and the parser return code for the command. This feature includes a mechanism for asynchronous notification to registered applications whenever the configuration changes. You can choose to have the notifications sent to the syslog.

For more information, see the *Configuration Change Notification and Logging* feature module:
http://www.cisco.com/en/US/docs/ios/fundamentals/configuration/guide/cf_config-logger_ps6350_TS_D_Products_Configuration_Guide_Chapter.html



Note Only CLI or HTTP changes are logged.

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 1-5](#) (optional)
- [Recalling Commands, page 1-5](#) (optional)
- [Disabling the Command History Feature, page 1-6](#) (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 1-4](#). These actions are optional.

Table 1-4 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.

Table 1-4 Recalling Commands (continued)

Action ¹	Result
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 appear is controlled by the setting of the terminal history global configuration command and the 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 1-6](#) (optional)
- [Editing Commands through Keystrokes, page 1-7](#) (optional)
- [Editing Command Lines that Wrap, page 1-8](#) (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 1-5 shows the keystrokes that you need to edit command lines. These keystrokes are optional.

Table 1-5 *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.
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-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.
	Press Ctrl-Y .	Recall the most recent entry in the buffer.
Delete entries if you make a mistake or change your mind.	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.
	Press the Delete or Backspace key.	Erase the character to the left of the cursor.
Capitalize or lowercase words or capitalize a set of letters.	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.
Designate a particular keystroke as an executable command, perhaps as a shortcut.	Press Esc C .	Capitalize at the cursor.
	Press Esc L .	Change the word at the cursor to lowercase.
	Press Esc U .	Capitalize letters from the cursor to the end of the word.
Designate a particular keystroke as an executable command, perhaps as a shortcut.	Press Ctrl-V or Esc Q .	

Table 1-5 Editing Commands through Keystrokes (continued)

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

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

Editing Command Lines that Wrap

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

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 1-7.

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* appear.

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
GigabitEthernet0/1 is up, line protocol is down
GigabitEthernet0/2 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.

You manage the switch stack and the stack member interfaces through the stack master. You cannot manage stack members on an individual switch basis. You can connect to the stack master through the console port of one or more stack members. Be careful with using multiple CLI sessions to the stack master. Commands you enter in one session are not displayed in the other sessions. Therefore, it is possible to lose track of the session from which you entered commands.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.



Note

We recommend using one CLI session when managing the switch stack.

If you want to configure a specific stack member port, you must include the stack member number in the CLI command interface notation. For more information about interface notations, see the [“Using Interface Configuration Mode” section on page 1-15](#).

To debug a specific stack member, you can access it from the stack master by using the **session stack-member-number** privileged EXEC command. The stack member number is appended to the system prompt. For example, `Switch-2#` is the prompt in privileged EXEC mode for stack member 2, and where the system prompt for the stack master is `Switch`. Only the **show** and **debug** commands are available in a CLI session to a specific stack member.

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 getting started guide that shipped with your switch. Then, to understand the boot process and the options available for assigning IP information, see [Chapter 1, “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 1-6.](#)

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, see the switch getting started guide or 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 1-6.](#) 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 1-41.](#) 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.



CHAPTER 1

Assigning the Switch IP Address and Default Gateway

This chapter describes how to create the initial switch configuration (for example, assigning the IP address and default gateway information) for the Catalyst 2960 or 2960-S switch by using a variety of automatic and manual methods. It also describes how to modify the switch startup configuration. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches.



Note

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release and the *Cisco IOS IP Command Reference, Volume 1 of 3: Addressing and Services* from Cisco.com page.

- [Understanding the Boot Process, page 1-1](#)
- [Assigning Switch Information, page 1-2](#)
- [Checking and Saving the Running Configuration, page 1-15](#)
- [Modifying the Startup Configuration, page 1-17](#)
- [Scheduling a Reload of the Software Image, page 1-22](#)

Understanding the Boot Process

To start your switch, you need to follow the procedures in the *Getting Started Guide* or the hardware installation guide for installing and powering on the switch and for setting up the initial switch configuration (IP address, subnet mask, default gateway, secret and Telnet passwords, and so forth).

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.
- Loads a default operating system software image into memory and boots up 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 a Software Failure”](#) section on page 1-2 and the [“Recovering from a Lost or Forgotten Password”](#) section on page 1-3.

**Note**

You can disable password recovery. For more information, see the [“Disabling Password Recovery”](#) section on page 1-5.

Before you can assign switch information, 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.

**Note**

If the data bits option is set to 8, set the parity option to none.

- Stop bits default is 1.
- Parity settings default is none.

Assigning Switch Information

You can assign IP information through the switch setup program, through a DHCP server, or manually.

Use the switch setup program if you want to be prompted for specific IP information. With this program, you can also configure a hostname and an enable secret password. It gives you the option of assigning a Telnet password (to provide security during remote management) and configuring your switch as a command or member switch of a cluster or as a standalone switch. For more information about the setup program, see the hardware installation guide.

The switch stack is managed through a single IP address. The IP address is a system-level setting and is not specific to the stack master or to any other stack member. You can still manage the stack through the same IP address even if you remove the stack master or any other stack member from the stack, provided there is IP connectivity.

Stacking is supported only on Catalyst 2960-S switches. Use a DHCP server for centralized control and automatic assignment of IP information after 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.

If you are an experienced user familiar with the switch configuration steps, manually configure the switch. Otherwise, use the setup program described previously.

- [Default Switch Information, page 1-3](#)
- [Understanding DHCP-Based Autoconfiguration, page 1-3](#)
- [Manually Assigning IP Information, page 1-14](#)

Default Switch Information

Table 1-1 shows the default switch information.

Table 1-1 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.
Hostname	The factory-assigned default hostname is <i>Switch</i> .
Telnet password	No password is defined.
Cluster command switch functionality	Disabled.
Cluster name	No cluster name is defined.

Understanding DHCP-Based Autoconfiguration

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. The switch can act as both a DHCP client and a DHCP server.

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 Trivial File Transfer Protocol (TFTP) server and a Domain Name System (DNS) server.

The DHCP server for your switch 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 device between your switch and the DHCP server. 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 up your switch, the DHCP client is invoked and requests configuration information from a DHCP server when the configuration file is not present on the switch. If the configuration file is present and the configuration includes the **ip address dhcp** interface configuration command on specific routed interfaces, the DHCP client is invoked and requests the IP address information for those interfaces.

Figure 1-1 shows the sequence of messages that are exchanged between the DHCP client and the DHCP server.

Figure 1-1 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 TFTP Server](#)” section on page 1-7.

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 switch. 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.

The DHCP hostname option allows a group of switches to obtain hostnames and a standard configuration from the central management DHCP server. A client (switch) includes in its DHCPDISCOVER message an option 12 field used to request a hostname and other configuration parameters from the DHCP server. The configuration files on all clients are identical except for their DHCP-obtained hostnames.

If a client has a default hostname (the **hostname name** global configuration command is not configured or the **no hostname** global configuration command is entered to remove the hostname), the DHCP hostname option is not included in the packet when you enter the **ip address dhcp** interface

configuration command. In this case, if the client receives the DHCP hostname option from the DHCP interaction while acquiring an IP address for an interface, the client accepts the DHCP hostname option and sets the flag to show that the system now has a hostname configured.

Understanding DHCP-based Autoconfiguration and Image Update

You can use the DHCP image upgrade features to configure a DHCP server to download both a new image and a new configuration file to one or more switches in a network. This helps ensure that each new switch added to a network receives the same image and configuration.

There are two types of DHCP image upgrades: DHCP autoconfiguration and DHCP auto-image update.

DHCP Autoconfiguration

DHCP autoconfiguration downloads a configuration file to one or more switches in your network from a DHCP server. The downloaded configuration file becomes the running configuration of the switch. It does not over write the bootup configuration saved in the flash, until you reload the switch.

DHCP Auto-Image Update

You can use DHCP auto-image upgrade with DHCP autoconfiguration to download both a configuration *and* a new image to one or more switches in your network. The switch (or switches) downloading the new configuration and the new image can be blank (or only have a default factory configuration loaded).

If the new configuration is downloaded to a switch that already has a configuration, the downloaded configuration is appended to the configuration file stored on the switch. (Any existing configuration is not overwritten by the downloaded one.)



Note

To enable a DHCP auto-image update on the switch, the TFTP server where the image and configuration files are located must be configured with the correct option 67 (the configuration filename), option 66 (the DHCP server hostname) option 150 (the TFTP server address), and option 125 (description of the file) settings.

For procedures to configure the switch as a DHCP server, see the [“Configuring DHCP-Based Autoconfiguration”](#) section on page 1-6 and the “Configuring DHCP” section of the “IP addressing and Services” section of the *Cisco IOS IP Configuration Guide, Release 12.2*.

After you install the switch in your network, the auto-image update feature starts. The downloaded configuration file is saved in the running configuration of the switch, and the new image is downloaded and installed on the switch. When you reboot the switch, the configuration is stored in the saved configuration on the switch.

Limitations and Restrictions

- The DHCP-based autoconfiguration with a saved configuration process stops if there is not at least one Layer 3 interface in an up state without an assigned IP address in the network.
- Unless you configure a timeout, the DHCP-based autoconfiguration with a saved configuration feature tries indefinitely to download an IP address.

- The auto-install process stops if a configuration file cannot be downloaded or if the configuration file is corrupted.

**Note**

The configuration file that is downloaded from TFTP is merged with the existing configuration in the running configuration but is not saved in the NVRAM unless you enter the **write memory** or **copy running-configuration startup-configuration** privileged EXEC command. Note that if the downloaded configuration is saved to the startup configuration, the feature is not triggered during subsequent system restarts.

Configuring DHCP-Based Autoconfiguration

- [DHCP Server Configuration Guidelines, page 1-6](#)
- [Configuring the TFTP Server, page 1-7](#)
- [Configuring the DNS, page 1-8](#)
- [Configuring the Relay Device, page 1-8](#)
- [Obtaining Configuration Files, page 1-9](#)
- [Example Configuration, page 1-9](#)

DHCP Server Configuration Guidelines

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

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)
- Hostname (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 previously, it replies to client requests with only those parameters that are configured. If the IP address and the subnet mask are not in the reply, the switch is not configured. If the router IP address or the 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 switch can act as a DHCP server. By default, the Cisco IOS DHCP server and relay agent features are enabled on your switch but are not configured. These features are not operational. If your DHCP server is a Cisco device, for additional information about configuring DHCP, see the “Configuring DHCP” section of the “IP Addressing and Services” section of the *Cisco IOS IP Configuration Guide* on Cisco.com page.

DHCP Server and Switch Stacks



Note

Stacking is supported only on Catalyst 2960-S switches.

The DHCP binding database is managed on the stack master. When a new stack master is assigned, the new master downloads the saved binding database from the TFTP server. If the stack master fails, all unsaved bindings are lost. The IP addresses associated with the lost bindings are released. You should configure an automatic backup by using the **ip dhcp database url [timeout seconds | write-delay seconds]** global configuration command.

When a stack merge occurs, the stack master that becomes a stack member loses all of the DHCP lease bindings. With a stack partition, the new master in the partition acts as a new DHCP server without any of the existing DHCP lease bindings.

For more information about the switch stack, see [Chapter 1, “Managing Switch Stacks.”](#)

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 previously), 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 1-8](#). 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, also referred to as a *relay agent*, when a switch sends broadcast packets that require a response from 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, 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 1-2](#), 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
```

Figure 1-2 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, it 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, it 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 `cisconet.cfg` default configuration file. (If the `network-config` file cannot be read, the switch reads the `cisconet.cfg` file.)

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

After obtaining its hostname from the default configuration file or the DHCP reply, the switch reads the configuration file that has the same name as its hostname (`hostname-config` or `hostname.cfg`, depending on whether `network-config` 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-config`, `cisconet.cfg`, or the hostname file, it reads the `router-config` file. If the switch cannot read the `router-config` file, it reads the `ciscotr.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 1-3 shows a sample network for retrieving IP information by using DHCP-based autoconfiguration.

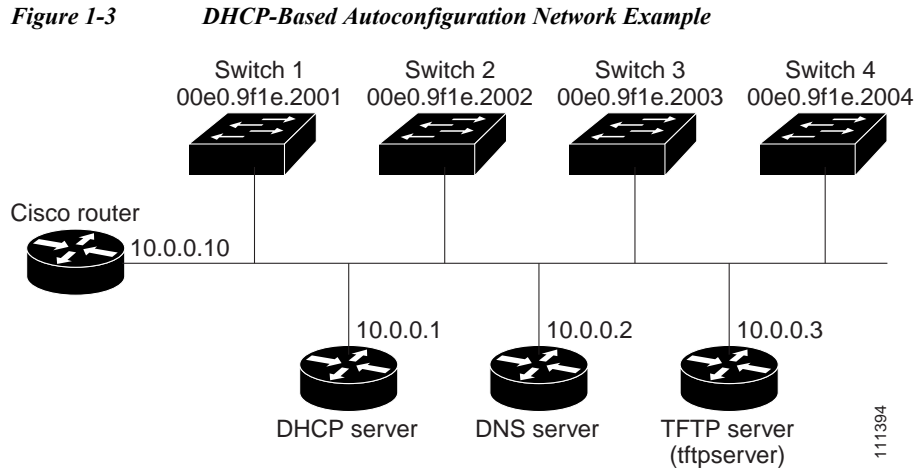


Table 1-2 shows the configuration of the reserved leases on the DHCP server.

Table 1-2 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	<i>tftpserver</i> or <i>10.0.0.3</i>	<i>tftpserver</i> or <i>10.0.0.3</i>	<i>tftpserver</i> or <i>10.0.0.3</i>	<i>tftpserver</i> or <i>10.0.0.3</i>
Boot filename (configuration file) (optional)	switcha-confg	switchb-confg	switchc-confg	switchd-confg
Hostname (optional)	switcha	switchb	switchc	switchd

DNS Server Configuration

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

```
prompt> cd /tftpserver/work/
prompt> ls
network-confg
switcha-confg
switchb-confg
switchc-confg
switchd-confg
prompt> cat network-confg
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 1-3](#), 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 hostname (switcha).
- It reads the configuration file that corresponds to its hostname; for example, it reads *switch1-config* from the TFTP server.

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

Configuring the DHCP Auto Configuration and Image Update Features

Using DHCP to download a new image and a new configuration to a switch requires that you configure at least two switches: One switch acts as a DHCP and TFTP server. The client switch is configured to download either a new configuration file or a new configuration file *and* a new image file.

Configuring DHCP Autoconfiguration (Only Configuration File)

Beginning in privileged EXEC mode, follow these steps to configure DHCP autoconfiguration of the TFTP and DHCP settings on a new switch to download a new configuration file.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip dhcp poolname	Create a name for the DHCP Server address pool, and enter DHCP pool configuration mode.
Step 3	bootfile filename	Specify the name of the configuration file that is used as a boot image.
Step 4	network network-number mask prefix-length	Specify the subnet network number and mask of the DHCP address pool. Note The prefix length specifies the number of bits that comprise the address prefix. The prefix is an alternative way of specifying the network mask of the client. The prefix length must be preceded by a forward slash (/).
Step 5	default-router address	Specify the IP address of the default router for a DHCP client.
Step 6	option 150 address	Specify the IP address of the TFTP server.

	Command	Purpose
Step 7	exit	Return to global configuration mode.
Step 8	tftp-server flash:filename.text	Specify the configuration file on the TFTP server.
Step 9	interface interface-id	Specify the address of the client that will receive the configuration file.
Step 10	no switchport	Put the interface into Layer 3 mode.
Step 11	ip address address mask	Specify the IP address and mask for the interface.
Step 12	end	Return to privileged EXEC mode.
Step 13	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to configure a switch as a DHCP server so that it will download a configuration file:

```
Switch# configure terminal
Switch(config)# ip dhcp pool pool1
Switch(dhcp-config)# network 10.10.10.0 255.255.255.0
Switch(dhcp-config)# bootfile config-boot.text
Switch(dhcp-config)# default-router 10.10.10.1
Switch(dhcp-config)# option 150 10.10.10.1
Switch(dhcp-config)# exit
Switch(config)# tftp-server flash:config-boot.text
Switch(config)# interface gigabitethernet1/0/4
Switch(config-if)# no switchport
Switch(config-if)# ip address 10.10.10.1 255.255.255.0
Switch(config-if)# end
```

Configuring DHCP Auto-Image Update (Configuration File and Image)

Beginning in privileged EXEC mode, follow these steps to configure DHCP autoconfiguration to configure TFTP and DHCP settings on a new switch to download a new image and a new configuration file.



Note

Before following the steps in this table, you must create a text file (for example, `autoinstall_dhcp`) that will be uploaded to the switch. In the text file, put the name of the image that you want to download. This image must be a tar and not a bin file.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip dhcp pool name	Create a name for the DHCP server address pool and enter DHCP pool configuration mode.
Step 3	bootfile filename	Specify the name of the file that is used as a boot image.
Step 4	network network-number mask prefix-length	Specify the subnet network number and mask of the DHCP address pool. Note The prefix length specifies the number of bits that comprise the address prefix. The prefix is an alternative way of specifying the network mask of the client. The prefix length must be preceded by a forward slash (/).
Step 5	default-router address	Specify the IP address of the default router for a DHCP client.

	Command	Purpose
Step 6	option 150 <i>address</i>	Specify the IP address of the TFTP server.
Step 7	option 125 <i>hex</i>	Specify the path to the text file that describes the path to the image file.
Step 8	copy tftp flash <i>filename.txt</i>	Upload the text file to the switch.
Step 9	copy tftp flash <i>imagename.tar</i>	Upload the tar file for the new image to the switch.
Step 10	exit	Return to global configuration mode.
Step 11	tftp-server flash: <i>config.text</i>	Specify the Cisco IOS configuration file on the TFTP server.
Step 12	tftp-server flash: <i>imagename.tar</i>	Specify the image name on the TFTP server.
Step 13	tftp-server flash: <i>filename.txt</i>	Specify the text file that contains the name of the image file to download
Step 14	interface <i>interface-id</i>	Specify the address of the client that will receive the configuration file.
Step 15	no switchport	Put the interface into Layer 3 mode.
Step 16	ip address <i>address mask</i>	Specify the IP address and mask for the interface.
Step 17	end	Return to privileged EXEC mode.
Step 18	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to configure a switch as a DHCP server so it downloads a configuration file:

```
Switch# config terminal
Switch(config)# ip dhcp pool pool1
Switch(dhcp-config)# network 10.10.10.0 255.255.255.0
Switch(dhcp-config)# bootfile config-boot.text
Switch(dhcp-config)# default-router 10.10.10.1
Switch(dhcp-config)# option 150 10.10.10.1
Switch(dhcp-config)# option 125 hex
0000.0009.0a05.08661.7574.6f69.6e73.7461.6c6c.5f64.686370
Switch(dhcp-config)# exit
Switch(config)# tftp-server flash:config-boot.text
Switch(config)# tftp-server flash:c2960 or 2960-S-lanbase-tar.122-46.SE.tar
Switch(config)# tftp-server flash:boot-config.text
Switch(config)# tftp-server flash:autoinstall_dhcp
Switch(config)# interface gigabitethernet1/0/4
Switch(config-if)# no switchport
Switch(config-if)# ip address 10.10.10.1 255.255.255.0
Switch(config-if)# end
```

Configuring the Client

Beginning in privileged EXEC mode, follow these steps to configure a switch to download a configuration file and new image from a DHCP server:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	boot host dhcp	Enable autoconfiguration with a saved configuration.
Step 3	boot host retry timeout <i>timeout-value</i>	(Optional) Set the amount of time the system tries to download a configuration file.
		Note If you do not set a timeout the system will indefinitely try to obtain an IP address from the DHCP server.

	Command	Purpose
Step 4	banner config-save ^C <i>warning-message</i> ^C	(Optional) Create warning messages to be displayed when you try to save the configuration file to NVRAM.
Step 5	end	Return to privileged EXEC mode.
Step 6	show boot	Verify the configuration.

This example uses a Layer 3 SVI interface on VLAN 99 to enable DHCP-based autoconfiguration with a saved configuration:

```
Switch# configure terminal
Switch(conf)# boot host dhcp
Switch(conf)# boot host retry timeout 300
Switch(conf)# banner config-save ^C Caution - Saving Configuration File to NVRAM May Cause
You to No longer Automatically Download Configuration Files at Reboot^C
Switch(config)# vlan 99
Switch(config-vlan)# interface vlan 99
Switch(config-if)# no shutdown
Switch(config-if)# end
Switch# show boot
BOOT path-list:
Config file:          flash:/config.text
Private Config file:  flash:/private-config.text
Enable Break:        no
Manual Boot:          no
HELPER path-list:
NVRAM/Config file
    buffer size:      32768
Timeout for Config
    Download:         300 seconds
Config Download
    via DHCP:         enabled (next boot: enabled)
Switch#
```



Note You should only configure and enable the Layer 3 interface. Do not assign an IP address or DHCP-based autoconfiguration with a saved configuration.

Manually Assigning IP Information

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

	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 VLAN range is 1 to 4094.
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 interfaces vlan <i>vlan-id</i>	Verify the configured IP address.
Step 8	show ip redirects	Verify the configured default gateway.
Step 9	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 1, “Administering the Switch.”](#)

Checking and Saving the Running Configuration

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

```
Switch# show running-config
Building configuration...

Current configuration: 1363 bytes
!
version 12.2
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname Switch A
!
enable secret 5 $1$ej9.$DMUvAUnZOAmvmgqBEzIxE0
!
.
<output truncated>
.
interface gigabitethernet6/0/1
ip address 172.20.137.50 255.255.255.0
!
interface gigabitethernet6/0/2
mvr type source

<output truncated>

...!
interface VLAN1
```

```

ip address 172.20.137.50 255.255.255.0
no ip directed-broadcast
!
ip default-gateway 172.20.137.1 !
!
snmp-server community private RW
snmp-server community public RO
snmp-server community private@es0 RW
snmp-server community public@es0 RO
snmp-server chassis-id 0x12
!
end

```

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

```

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

```

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

For more information about alternative locations from which to copy the configuration file, see [Appendix 1, “Working with the Cisco IOS File System, Configuration Files, and Software Images.”](#)

Configuring the NVRAM Buffer Size

The default NVRAM buffer size is 512 KB. In some cases, the configuration file might be too large to save to NVRAM. Typically, this occurs when you have many switches in a switch stack. You can configure the size of the NVRAM buffer to support larger configuration files. The new NVRAM buffer size is synced to all current and new member switches.



Note

After you configure the NVRAM buffer size, reload the switch or switch stack.

When you add a switch to a stack and the NVRAM size differs, the new switch syncs with the stack and reloads automatically.

Beginning in privileged EXEC mode, follow these steps to configure the NVRAM buffer size:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	boot buffersize <i>size</i>	Configure the NVRAM buffersize in KB. The valid range for <i>size</i> is from 4096 to 1048576.
Step 3	end	Return to privileged EXEC mode.
Step 4	show boot	Verify the configuration.

This example shows how to configure the NVRAM buffer size:

```
Switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# boot buffersize 524288
Switch(config)# end
Switch# show boot
BOOT path-list      :
Config file         : flash:/config.text
Private Config file : flash:/private-config.text
Enable Break       : no
Manual Boot        : no
HELPER path-list   :
Auto upgrade       : yes
Auto upgrade path  :
NVRAM/Config file
  buffer size:     524288
Timeout for Config
  Download:        300 seconds
Config Download
  via DHCP:        enabled (next boot: enabled)
Switch#
```

Modifying the Startup Configuration

These sections describe how to modify the switch startup configuration:

- [Default Boot Configuration, page 1-18](#)
- [Automatically Downloading a Configuration File, page 1-18](#)
- [Booting Manually, page 1-19](#)
- [Booting a Specific Software Image, page 1-19](#)
- [Controlling Environment Variables, page 1-20](#)

See also “Stack Configuration Files” section on page 1-13 and Appendix 1, “Working with the Cisco IOS File System, Configuration Files, and Software Images,” for information about switch stackconfiguration files.



Note

Stacking is supported only on Catalyst 2960-S switches.

Default Boot Configuration

Table 1-3 Default Boot Configuration

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

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 1-3.

Specifying the Filename to Read and Write the System Configuration

By default, the Cisco 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 will be loaded during the next boot-up cycle.



Note This command only works properly from a standalone switch.

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	<p>Specify the configuration file to load during the next boot-up cycle.</p> <p>For <i>file-url</i>, specify the path (directory) and the configuration filename.</p> <p>Filenames and directory names are case sensitive.</p>
Step 3	end	Return to privileged EXEC mode.

	Command	Purpose
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 up; however, you can configure it to manually boot up.



Note This command only works properly from a standalone switch.

Beginning in privileged EXEC mode, follow these steps to configure the switch to manually boot up 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 up 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 MANUAL_BOOT 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 up the system, use the boot filesystem:/file-url boot loader command. <ul style="list-style-type: none"> For <i>filesystem:</i>, use flash: for the system board flash device. For <i>file-url</i>, specify the path (directory) and the name of the bootable image. Filenames and directory names are case sensitive.
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 up 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 up.

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 <i>filesystem:/file-url</i>	Configure the switch to boot a specific image in flash memory during the next boot cycle. <ul style="list-style-type: none"> For <i>filesystem:</i>, use flash: for the system board flash device. For <i>file-url</i>, specify the path (directory) and the name of the bootable image. <p>If you enter this command on a stack master, the specified software image is loaded only on the stack master during the next boot cycle.</p> <p>Note Stacking is supported only on Catalyst 2960-S switches.</p> <p>Filenames and directory names are case sensitive.</p>
Step 3	boot system switch { <i>number</i> all }	(Optional) Specify the switch members on which the system image is loaded during the next boot cycle: <ul style="list-style-type: none"> Use <i>number</i> to specify a stack member. (Specify only one stack member.) Use all to specify all stack members.
Step 4	end	Return to privileged EXEC mode.
Step 5	show boot	Verify your entries. <p>The boot system global command changes the setting of the BOOT environment variable.</p> <p>During the next boot cycle, the switch attempts to automatically boot up the system using information in the BOOT environment variable.</p>
Step 6	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 b/s. 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 appears.

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 Cisco IOS configuration file. For example, the name of a boot loader helper file, which extends or patches the functionality of the boot loader can be stored as an environment variable.
- Data that controls code, which is responsible for reading the Cisco IOS configuration file. For example, the name of the Cisco IOS configuration file can be stored as an environment variable.

You can change the settings of the environment variables by accessing the boot loader or by using Cisco 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, see the command reference for this release.

Table 1-4 describes the function of the most common environment variables.

Table 1-4 Environment Variables

Variable	Boot Loader Command	Cisco IOS Global Configuration Command
BOOT	<p>set BOOT <i>filesystem:/file-url ...</i></p> <p>A semicolon-separated list of executable files to try to load and execute when automatically booting. If the BOOT environment variable is not set, the system attempts to load and execute the first executable image it can find by using a recursive, depth-first search through the flash file system. If the BOOT variable is set but the specified images cannot be loaded, the system attempts to boot the first bootable file that it can find in the flash file system.</p>	<p>boot system {<i>filesystem:/file-url ...</i> switch {<i>number</i> all}}</p> <p>Specifies the Cisco IOS image to load during the next boot cycle and the stack members on which the image is loaded. This command changes the setting of the BOOT environment variable.</p> <p>Note Stacking is not supported on Catalyst 2960 switches.</p>
MANUAL_BOOT	<p>set MANUAL_BOOT yes</p> <p>Decides whether the switch automatically or manually boots up.</p> <p>Valid values are 1, yes, 0, and no. If it is set to no or 0, the boot loader attempts to automatically boot up the system. If it is set to anything else, you must manually boot up the switch from the boot loader mode.</p>	<p>boot manual</p> <p>Enables manually booting up the switch during the next boot cycle and changes the setting of the MANUAL_BOOT environment variable.</p> <p>The next time you reboot the system, the switch is in boot loader mode. To boot up the system, use the boot flash:<i>filesystem:/file-url</i> boot loader command, and specify the name of the bootable image.</p>
CONFIG_FILE	<p>set CONFIG_FILE flash:<i>/file-url</i></p> <p>Changes the filename that Cisco IOS uses to read and write a nonvolatile copy of the system configuration.</p>	<p>boot config-file flash:<i>/file-url</i></p> <p>Specifies the filename that Cisco IOS uses to read and write a nonvolatile copy of the system configuration. This command changes the CONFIG_FILE environment variable.</p>

Table 1-4 Environment Variables (continued)

Variable	Boot Loader Command	Cisco IOS Global Configuration Command
SWITCH_NUMBER	set SWITCH_NUMBER <i>stack-member-number</i> Changes the member number of a stack member.	switch <i>current-stack-member-number</i> renumber <i>new-stack-member-number</i> Note Stacking is not supported on Catalyst 2960 switches. Changes the member number of a stack member.
SWITCH_PRIORITY	set SWITCH_PRIORITY <i>stack-member-number</i> Changes the priority value of a stack member.	switch <i>stack-member-number</i> priority <i>priority-number</i> Note Stacking is not supported on Catalyst 2960 switches. Changes the priority value of a stack member.

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.

To reload a specific switch in a switch stack, use the **reload slot** *stack-member-number* privileged EXEC command.



Note

Stacking is supported only on Catalyst 2960-S switches.

- **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 up, it reboots itself. Use the **reload** command after you save the switch configuration information to the startup configuration (**copy running-config startup-config**).

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

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

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

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

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

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

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

Displaying Scheduled Reload Information

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

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



CHAPTER 1

Configuring Cisco IOS Configuration Engine

This chapter describes how to configure the feature on the Catalyst 2960 and 2960-S switches.



Note

For complete configuration information for the Cisco Configuration Engine, go to http://www.cisco.com/en/US/products/sw/netmgmtsw/ps4617/tsd_products_support_series_home.html

For complete syntax and usage information for the commands used in this chapter, go to the *Cisco IOS Network Management Command Reference, Release 12.4*:

http://www.cisco.com/en/US/docs/ios/netmgmt/command/reference/nm_book.html

- [Understanding Cisco Configuration Engine Software, page 1-1](#)
- [Understanding Cisco IOS Agents, page 1-5](#)
- [Configuring Cisco IOS Agents, page 1-6](#)
- [Displaying CNS Configuration, page 1-13](#)

Understanding Cisco Configuration Engine Software

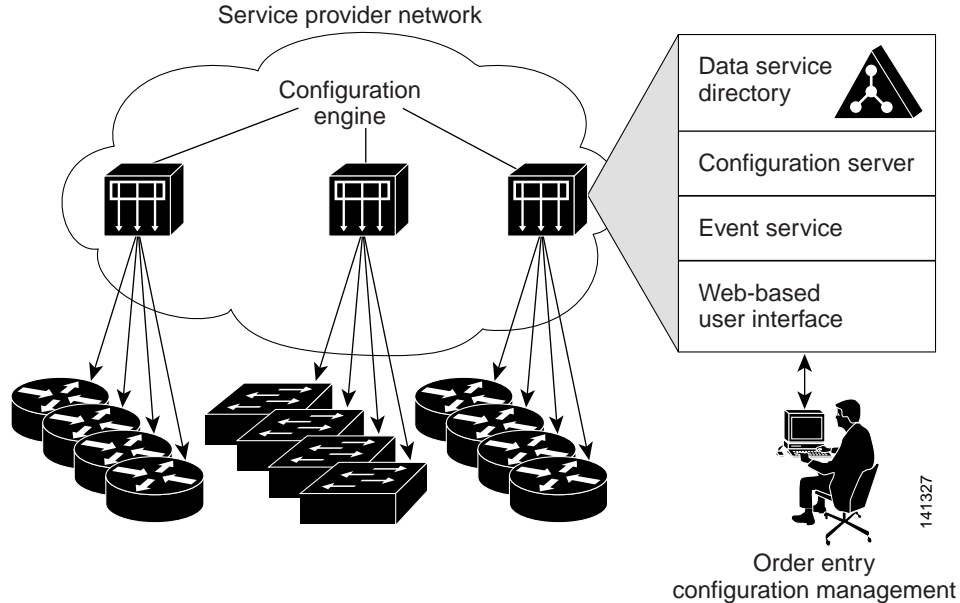
The Cisco Configuration Engine is network management software that acts as a configuration service for automating the deployment and management of network devices and services (see [Figure 1-1](#)). Each Configuration Engine manages a group of Cisco devices (switches and routers) and the services that they deliver, storing their configurations and delivering them as needed. The Configuration Engine 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 Engine 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 Engine supports an embedded Directory Service. In this mode, no external directory or other data store is required. In server mode, the Configuration Engine supports the use of a user-defined external directory.

Figure 1-1 Configuration Engine Architectural Overview



- [Configuration Service, page 1-2](#)
- [Event Service, page 1-3](#)
- [What You Should Know About the CNS IDs and Device Hostnames, page 1-3](#)

Configuration Service

The Configuration Service is the core component of the Cisco Configuration Engine. It consists of a configuration server that works with Cisco IOS CNS agents on the switch. The 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 Configuration Service when they start up on the network for the first time.

The 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 Cisco IOS agent can perform a syntax check on received configuration files and publish events to show 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.

Event Service

The Cisco Configuration Engine uses the Event Service for receipt and generation of configuration events. The event agent is on the switch and facilitates the communication between the switch and the event gateway on the Configuration Engine.

The Event Service is a highly capable publish-and-subscribe communication method. The 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 Engine includes the NameSpace Mapper (NSM) that provides a lookup service for managing logical groups of devices based on application, device 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 changes your event subject-name strings to those known by Cisco 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.

What You Should Know About the CNS IDs and Device Hostnames

The Configuration Engine 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 Engine 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 Engine uses both the event bus and the configuration server to provide configurations to devices, you must define both ConfigID and Device ID 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 Engine 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 Engine.

The ConfigID is fixed at startup time and cannot be changed until the device restarts, even if the switch hostname 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 Engine.

The origin of the DeviceID is defined by the Cisco IOS hostname of the switch. However, the DeviceID variable and its usage reside within the event gateway 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 hostname to the event gateway immediately after the successful connection to the event gateway. The event gateway couples the DeviceID value to the Cisco IOS hostname each time this connection is established. The event gateway caches this DeviceID value for the duration of its connection to the switch.

Hostname and DeviceID

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

When changing the switch hostname 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 hostname to the event gateway. The event gateway redefines the DeviceID to the new value.



Caution

When using the Configuration Engine user interface, you must first set the DeviceID field to the hostname 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 Hostname, DeviceID, and ConfigID

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

In server mode, the hostname 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 Engine.



Note

For more information about running the setup program on the Configuration Engine, see the Configuration Engine setup and configuration guide:

http://www.cisco.com/en/US/products/sw/netmgtsw/ps4617/prod_installation_guides_list.html

Understanding Cisco IOS Agents

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

- [Initial Configuration, page 1-5](#)
- [Incremental \(Partial\) Configuration, page 1-6](#)
- [Synchronized Configuration, page 1-6](#)

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 Cisco IOS agents initiate communication with the Configuration Engine by using the appropriate ConfigID and EventID. The Configuration Engine maps the Config ID to a template and downloads the full configuration file to the switch.

[Figure 1-2](#) shows a sample network configuration for retrieving the initial bootstrap configuration file by using DHCP-based autoconfiguration.

Figure 1-2 Initial Configuration Overview



Incremental (Partial) Configuration

After the network is running, new services can be added by using the Cisco IOS agent. Incremental (partial) configurations can be sent to the switch. The actual configuration can be sent 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 Cisco IOS Agents

The Cisco IOS 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 1-6](#). If you want to change the configuration or install a custom configuration, see these sections for instructions:

- [Enabling the CNS Event Agent, page 1-7](#)
- [Enabling the Cisco IOS CNS Agent, page 1-9](#)

Enabling Automated CNS Configuration

To enable automated CNS configuration of the switch, you must first complete the prerequisites in [Table 1-1](#). 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 1-5](#). When the full configuration file is loaded on your switch, you need to do nothing else.

Table 1-1 Prerequisites for Enabling Automatic Configuration

Device	Required Configuration
Access switch	Factory default (no configuration file)
Distribution switch	<ul style="list-style-type: none"> • IP helper address • Enable DHCP relay agent • IP routing (if used as default gateway)

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

Device	Required Configuration
DHCP server	<ul style="list-style-type: none"> • IP address assignment • TFTP server IP address • Path to bootstrap configuration file on the TFTP server • Default gateway IP address
TFTP server	<ul style="list-style-type: none"> • A bootstrap configuration file that includes the CNS configuration commands that enable the switch to communicate with the Configuration Engine • The switch configured to use either the switch MAC address or the serial number (instead of the default hostname) to generate the ConfigID and EventID • The CNS event agent configured to push the configuration file to the switch
CNS Configuration Engine	One or more templates for each type of device, with the ConfigID of the device mapped to the template.

**Note**

For more information about running the setup program and creating templates on the Configuration Engine, see the *Cisco Configuration Engine Installation and Setup Guide, 1.5 for Linux*: http://www.cisco.com/en/US/docs/net_mgmt/configuration_engine/1.5/installation_linux/guide/setup_1.html

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>hostname</i> <i>ip-address</i> } [<i>port-number</i>] [backup] [failover-time <i>seconds</i>] [keepalive <i>seconds</i> <i>retry-count</i>] [reconnect <i>time</i>] [source <i>ip-address</i>]	<p>Enable the event agent, and enter the gateway parameters.</p> <ul style="list-style-type: none"> For {<i>hostname</i> <i>ip-address</i>}, enter either the hostname or the IP address 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 failover-time <i>seconds</i>, enter how long the switch waits for the primary gateway route after the route to the backup gateway is established. (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 reconnect <i>time</i>, enter the maximum time interval that the switch waits before trying to reconnect to the event gateway. (Optional) For source <i>ip-address</i>, enter the source IP address of this device. <p>Note Though visible in the command-line help string, the encrypt and the clock-timeout <i>time</i> keywords are not supported.</p>
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 Cisco IOS CNS Agent

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

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

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 template connect <i>name</i>	Enter CNS template connect configuration mode, and specify the name of the CNS connect template.
Step 3	cli <i>config-text</i>	Enter a command line for the CNS connect template. Repeat this step for each command line in the template.
Step 4		Repeat Steps 2 to 3 to configure another CNS connect template.
Step 5	exit	Return to global configuration mode.
Step 6	cns connect <i>name</i> [retries <i>number</i>] [retry-interval <i>seconds</i>] [sleep <i>seconds</i>] [timeout <i>seconds</i>]	<p>Enter CNS connect configuration mode, specify the name of the CNS connect profile, and define the profile parameters. The switch uses the CNS connect profile to connect to the Configuration Engine.</p> <ul style="list-style-type: none"> • Enter the name of the CNS connect profile. • (Optional) For retries <i>number</i>, enter the number of connection retries. The range is 1 to 30. The default is 3. • (Optional) For retry-interval <i>seconds</i>, enter the interval between successive connection attempts to the Configuration Engine. The range is 1 to 40 seconds. The default is 10 seconds. • (Optional) For sleep <i>seconds</i>, enter the amount of time before which the first connection attempt occurs. The range is 0 to 250 seconds. The default is 0. • (Optional) For timeout <i>seconds</i>, enter the amount of time after which the connection attempts end. The range is 10 to 2000 seconds. The default is 120.

	Command	Purpose
Step 7	discover { controller <i>controller-type</i> dcli [subinterface <i>subinterface-number</i>] interface [<i>interface-type</i>] line <i>line-type</i> }	Specify the interface parameters in the CNS connect profile. <ul style="list-style-type: none"> For controller <i>controller-type</i>, enter the controller type. For dcli, enter the active data-link connection identifiers (DLCIs). (Optional) For subinterface <i>subinterface-number</i>, specify the point-to-point subinterface number that is used to search for active DLCIs. For interface [<i>interface-type</i>], enter the type of interface. For line <i>line-type</i>, enter the line type.
Step 8	template <i>name</i> [... <i>name</i>]	Specify the list of CNS connect templates in the CNS connect profile to be applied to the switch configuration. You can specify more than one template.
Step 9		Repeat Steps 7 to 8 to specify more interface parameters and CNS connect templates in the CNS connect profile.
Step 10	exit	Return to global configuration mode.
Step 11	hostname <i>name</i>	Enter the hostname for the switch.
Step 12	ip route <i>network-number</i>	(Optional) Establish a static route to the Configuration Engine whose IP address is <i>network-number</i> .
Step 13	cns id <i>interface num</i> { dns-reverse ipaddress mac-address } [event] [image] or cns id { hardware-serial hostname string <i>string</i> udi } [event] [image]	(Optional) Set the unique EventID or ConfigID used by the Configuration Engine. <ul style="list-style-type: none"> 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 hostname 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. (Optional) Enter image to set the ID to be the image-id value used to identify the switch. <p>Note If both the event and image keywords are omitted, the image-id value is used to identify the switch.</p> <ul style="list-style-type: none"> For { hardware-serial hostname string <i>string</i> udi }, enter hardware-serial to set the switch serial number as the unique ID, enter hostname (the default) to select the switch hostname as the unique ID, enter an arbitrary text string for string <i>string</i> as the unique ID, or enter udi to set the unique device identifier (UDI) as the unique ID.

	Command	Purpose
Step 14	cns config initial { <i>hostname</i> <i>ip-address</i> } [<i>port-number</i>] [event] [no-persist] [page <i>page</i>] [source <i>ip-address</i>] [syntax-check]	<p>Enable the Cisco IOS agent, and initiate an initial configuration.</p> <ul style="list-style-type: none"> For {<i>hostname</i> <i>ip-address</i>}, enter the hostname or the IP address 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 <i>page</i>, enter the web page of the initial configuration. The default is /Config/config/asp. (Optional) Enter source <i>ip-address</i> to use for source IP address. (Optional) Enable syntax-check to check the syntax when this parameter is entered. <p>Note Though visible in the command-line help string, the encrypt, status url, and inventory keywords are not supported.</p>
Step 15	end	Return to privileged EXEC mode.
Step 16	show cns config connections	Verify information about the configuration agent.
Step 17	show running-config	Verify your entries.

To disable the CNS Cisco IOS 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 when the switch configuration is unknown (the CNS Zero Touch feature).

```
Switch(config)# cns template connect template-dhcp
Switch(config-tmpl-conn)# cli ip address dhcp
Switch(config-tmpl-conn)# exit
Switch(config)# cns template connect ip-route
Switch(config-tmpl-conn)# cli ip route 0.0.0.0 0.0.0.0 ${next-hop}
Switch(config-tmpl-conn)# exit
Switch(config)# cns connect dhcp
Switch(config-cns-conn)# discover interface gigabitethernet
Switch(config-cns-conn)# template template-dhcp
Switch(config-cns-conn)# template ip-route
Switch(config-cns-conn)# exit
Switch(config)# hostname RemoteSwitch
RemoteSwitch(config)# cns config initial 10.1.1.1 no-persist
```

This example shows how to configure an initial configuration on a remote switch when the switch IP address is known. The Configuration Engine IP address is 172.28.129.22.

```
Switch(config)# cns template connect template-dhcp
Switch(config-tmpl-conn)# cli ip address dhcp
Switch(config-tmpl-conn)# exit
Switch(config)# cns template connect ip-route
Switch(config-tmpl-conn)# cli ip route 0.0.0.0 0.0.0.0 ${next-hop}
Switch(config-tmpl-conn)# exit
Switch(config)# cns connect dhcp
Switch(config-cns-conn)# discover interface gigabitethernet
Switch(config-cns-conn)# template template-dhcp
Switch(config-cns-conn)# template ip-route
Switch(config-cns-conn)# exit
Switch(config)# hostname RemoteSwitch
RemoteSwitch(config)# ip route 172.28.129.22 255.255.255.255 11.11.11.1
RemoteSwitch(config)# cns id ethernet 0 ipaddress
RemoteSwitch(config)# cns config initial 172.28.129.22 no-persist
```

Enabling a Partial Configuration

Beginning in privileged EXEC mode, follow these steps to enable the Cisco IOS 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. <ul style="list-style-type: none"> For {<i>ip-address</i> <i>hostname</i>}, enter the IP address or the hostname 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. <p>Note Though visible in the command-line help string, the encrypt keyword is not supported.</p>
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 Cisco IOS 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

Table 1-2 Privileged EXEC show Commands

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



CHAPTER 1

Administering the Switch

This chapter describes how to perform one-time operations to administer the Catalyst 2960 and 2960-S switches. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

This chapter consists of these sections:

- [Identifying the Switch Image, page 1-1](#)
- [Managing the System Time and Date, page 1-2](#)
- [Configuring a System Name and Prompt, page 1-8](#)
- [Creating a Banner, page 1-11](#)
- [Managing the MAC Address Table, page 1-13](#)
- [Managing the ARP Table, page 1-24](#)

Identifying the Switch Image

The Catalyst 2960 and 2960-S switches run one of these images:

- The LAN base software image provides enterprise-class intelligent services such as access control lists (ACLs) and quality of service (QoS) features. On a Catalyst 2960-S switch, stacking is also supported.
- The LAN Lite image provides reduced functionality.

The Catalyst 2960-S ships with a universal image that includes cryptographic functionality. The software image on the switch is either the LAN base or LAN Lite image, depending on the switch model. To determine which image your switch is running:

- Switches running the LAN Lite image do not support the FlexStack module. They do not have a FlexStack module slot on the rear of the switch.
- On the front of the switch, the label in the top right corner ends in -L if the switch model runs the LAN base image and -S if the switch model runs the LAN Lite image.

- Enter the show version privileged EXEC command. The line that shows the product ID also ends in either -L (if running the LAN base image) or -S (if running the LAN Lite image). For example, WS-C2960S-48PD-L is running LAN base; WS-C2960S-24TS-S is running LAN Lite image.
- Enter the show license privileged EXEC command, and see which is the active image:

```
Switch# show license
Index 1 Feature: lanlite
      Period left: 0 minute 0 second
Index 2 Feature: lanbase
      Period left: Life time
      License Type: Permanent
      License State: Active, In Use
      License Priority: Medium
      License Count: Non-Counted
```

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, see the *Cisco IOS Configuration Fundamentals Command Reference* on Cisco.com.

These sections contain this configuration information:

- [Understanding the System Clock, page 1-2](#)
- [Understanding Network Time Protocol, page 1-3](#)
- [NTP Version 4, page 1-4](#)
- [Configuring Time and Date Manually, page 1-5](#)

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:

- NTP
- 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 appears correctly 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 1-5](#).

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 1-1 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 1-1 Typical NTP Network Configuration



101349

If the network is isolated from the Internet, Cisco's implementation of NTP allows a device to act as if it is synchronized through NTP, when in fact it has learned 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.

NTP Version 4

NTP version 4 is implemented on the switch. NTPv4 is an extension of NTP version 3. NTPv4 supports both IPv4 and IPv6 and is backward-compatible with NTPv3.

NTPv4 provides these capabilities:

- Support for IPv6.
- Improved security compared to NTPv3. The NTPv4 protocol provides a security framework based on public key cryptography and standard X509 certificates.
- Automatic calculation of the time-distribution hierarchy for a network. Using specific multicast groups, NTPv4 automatically configures the hierarchy of the servers to achieve the best time accuracy for the lowest bandwidth cost. This feature leverages site-local IPv6 multicast addresses.

For details about configuring NTPv4, see the “[Implementing NTPv4 in IPv6](#)” chapter of the *Cisco IOS IPv6 Configuration Guide, Release 12.4T*.

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.

These sections contain this configuration information:

- [Setting the System Clock, page 1-5](#)
- [Displaying the Time and Date Configuration, page 1-5](#)
- [Configuring the Time Zone, page 1-6](#)
- [Configuring Summer Time \(Daylight Saving Time\), page 1-7](#)

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	<code>clock set hh:mm:ss day month year</code> or <code>clock set hh:mm:ss month day year</code>	Manually set the system clock using one of these formats. <ul style="list-style-type: none"> • For <i>hh:mm:ss</i>, specify the time in hours (24-hour format), minutes, and seconds. The time specified is relative to the configured time zone. • For <i>day</i>, specify the day by date in the month. • For <i>month</i>, specify the month by name. • For <i>year</i>, specify the year (no abbreviation).

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. <ul style="list-style-type: none"> • For <i>zone</i>, enter the name of the time zone to be displayed when standard time is in effect. The default is UTC. • For <i>hours-offset</i>, enter the hours offset from UTC. • (Optional) For <i>minutes-offset</i>, enter the minutes offset from UTC.
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</i> <i>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. <ul style="list-style-type: none"> For <i>zone</i>, specify the name of the time zone (for example, PDT) to be displayed when summer time is in effect. (Optional) For <i>week</i>, specify the week of the month (1 to 5 or last). (Optional) For <i>day</i>, specify the day of the week (Sunday, Monday...). (Optional) For <i>month</i>, specify the month (January, February...). (Optional) For <i>hh:mm</i>, specify the time (24-hour format) in hours and minutes. (Optional) For <i>offset</i>, specify the number of minutes to add during summer time. The default is 60.
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</i> [<i>offset</i>]] or clock summer-time zone date [<i>date month year hh:mm date month year hh:mm</i> [<i>offset</i>]]	Configure summer time to start on the first date and end on the second date. Summer time is disabled by default. <ul style="list-style-type: none"> For <i>zone</i>, specify the name of the time zone (for example, PDT) to be displayed when summer time is in effect. (Optional) For <i>week</i>, specify the week of the month (1 to 5 or last). (Optional) For <i>day</i>, specify the day of the week (Sunday, Monday...). (Optional) For <i>month</i>, specify the month (January, February...). (Optional) For <i>hh:mm</i>, specify the time (24-hour format) in hours and minutes. (Optional) For <i>offset</i>, specify the number of minutes to add during summer time. The default is 60.
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, 2000, at 02:00, and end on April 26, 2001, at 02:00:

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

Configuring a System Name and Prompt

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

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

If you are accessing a stack member through the stack master, you must use the **session stack-member-number** privileged EXEC command. The stack member number range is from 1 through 9. When you use this command, the stack member number is appended to the system prompt. For example, *Switch-2#* is the prompt in privileged EXEC mode for stack member 2, and the system prompt for the switch stack is *Switch*.

For complete syntax and usage information for the commands used in this section, see the *Cisco IOS Configuration Fundamentals Command Reference* and the *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols*.

These sections contain this configuration information:

- [Default System Name and Prompt Configuration, page 1-9](#)
- [Configuring a System Name, page 1-9](#)
- [Understanding DNS, page 1-9](#)

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 name	Manually configure a system name. The default setting is <i>switch</i> . The name must follow the rules for ARPANET hostnames. 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.

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

Understanding DNS

The DNS protocol controls the Domain Name System (DNS), a distributed database with which you can map hostnames to IP addresses. When you configure DNS on your switch, you can substitute the hostname 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 File Transfer Protocol (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 hostnames, specify the name server that is present on your network, and enable the DNS.

These sections contain this configuration information:

- [Default DNS Configuration, page 1-10](#)
- [Setting Up DNS, page 1-10](#)
- [Displaying the DNS Configuration, page 1-11](#)

Default DNS Configuration

Table 1-1 shows the default DNS configuration.

Table 1-1 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 hostnames (names without a dotted-decimal domain name). Do not include the initial period that separates an unqualified name from the domain name. At boot-up time, no domain name is configured; however, if the switch configuration comes from a BOOTP or Dynamic Host Configuration Protocol (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 hostname-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).

	Command	Purpose
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 appears after the MOTD banner and before the login prompts.



Note

For complete syntax and usage information for the commands used in this section, see the *Cisco IOS Configuration Fundamentals Command Reference, Release 12.4* on Cisco.com.

These sections contain this configuration information:

- [Default Banner Configuration, page 1-11](#)
- [Configuring a Message-of-the-Day Login Banner, page 1-12](#)
- [Configuring a Login Banner, page 1-13](#)

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 that appears 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 c message c	Specify the login message. For <i>c</i> , enter the delimiting character of your choice, for example, a pound sign (#), and press the Return key. The delimiting character signifies the beginning and end of the banner text. Characters after the ending delimiter are discarded. For <i>message</i> , enter a login message up to 255 characters. You cannot use the delimiting character in the message.
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, see the command reference for this release.

These sections contain this configuration information:

- [Building the Address Table, page 1-14](#)
- [MAC Addresses and VLANs, page 1-14](#)
- [MAC Addresses and Switch Stacks, page 1-15](#)
- [Default MAC Address Table Configuration, page 1-15](#)
- [Changing the Address Aging Time, page 1-15](#)
- [Removing Dynamic Address Entries, page 1-16](#)
- [Configuring MAC Address Change Notification Traps, page 1-16](#)
- [Configuring MAC Address Move Notification Traps, page 1-18](#)
- [Configuring MAC Threshold Notification Traps, page 1-19](#)
- [Adding and Removing Static Address Entries, page 1-20](#)
- [Configuring Unicast MAC Address Filtering, page 1-21](#)
- [Disabling MAC Address Learning on a VLAN, page 1-22](#)
- [Displaying Address Table Entries, page 1-23](#)

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 on a standalone switch or on the switch stack. 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 be forwarded to port 1 in VLAN 1 and ports 9, 10, and 1 in VLAN 5.

Each VLAN maintains its own logical address table. A known address in one VLAN is unknown in another until it is learned or statically associated with a port in the other VLAN.

MAC Addresses and Switch Stacks

The MAC address tables on all stack members are synchronized. At any given time, each stack member has the same copy of the address tables for each VLAN. When an address ages out, the address is removed from the address tables on all stack members. When a switch joins a switch stack, that switch receives the addresses for each VLAN learned on the other stack members. When a stack member leaves the switch stack, the remaining stack members age out or remove all addresses learned by the former stack member.

Default MAC Address Table Configuration

Table 1-2 shows the default MAC address table configuration.

Table 1-2 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> , valid IDs are 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 Change Notification Traps

MAC address change notification tracks users on a network by storing the MAC address change activity. When the switch learns or removes a MAC address, an SNMP notification trap can be 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 to reduce network traffic. The MAC notification history table stores MAC address activity for each port for which the trap is set. MAC address change notifications are generated for dynamic and secure MAC addresses. Notifications 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 change 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 notification-type</i>	Specify the recipient of the trap message. <ul style="list-style-type: none"> For <i>host-addr</i>, specify the name or address of the NMS. Specify traps (the default) to send SNMP traps to the host. Specify informs to send SNMP informs to the host. Specify the SNMP version to support. Version 1, the default, is not available with informs. For <i>community-string</i>, specify the string to send with the notification operation. Though you can set this string by using the snmp-server host command, we recommend that you define this string by using the snmp-server community command before using the snmp-server host command. For <i>notification-type</i>, use the mac-notification keyword.
Step 3	snmp-server enable traps mac-notification change	Enable the switch to send MAC address change notification traps to the NMS.
Step 4	mac address-table notification change	Enable the MAC address change notification feature.

	Command	Purpose
Step 5	mac address-table notification change [<i>interval value</i>] [<i>history-size value</i>]	Enter the trap interval time and the history table size. <ul style="list-style-type: none"> (Optional) For interval value, 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 value, specify the maximum number of entries in the MAC notification history table. The range is 0 to 500; the default is 1.
Step 6	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface on which to enable the SNMP MAC address notification trap.
Step 7	snmp trap mac-notification change { added removed }	Enable the MAC address change notification trap on the interface. <ul style="list-style-type: none"> Enable the trap when a MAC address is added on this interface. Enable the trap when a MAC address is removed from this interface.
Step 8	end	Return to privileged EXEC mode.
Step 9	show mac address-table notification change interface show running-config	Verify your entries.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable MAC address-change notification traps, use the **no snmp-server enable traps mac-notification change** global configuration command. To disable the MAC address-change notification traps on a specific interface, use the **no snmp trap mac-notification change {added | removed}** interface configuration command. To disable the MAC address-change notification feature, use the **no mac address-table notification change** 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-change notification feature, set the interval time to 123 seconds, set the history-size to 100 entries, and enable traps whenever a MAC address is added on the specified port.

```
Switch(config)# snmp-server host 172.20.10.10 traps private mac-notification
Switch(config)# snmp-server enable traps mac-notification change
Switch(config)# mac address-table notification change
Switch(config)# mac address-table notification change interval 123
Switch(config)# mac address-table notification change history-size 100
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# snmp trap mac-notification change added
```

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

Configuring MAC Address Move Notification Traps

When you configure MAC-move notification, an SNMP notification is generated and sent to the network management system whenever a MAC address moves from one port to another within the same VLAN.

Beginning in privileged EXEC mode, follow these steps to configure the switch to send MAC address-move 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 notification-type</i>	Specify the recipient of the trap message. <ul style="list-style-type: none"> For <i>host-addr</i>, specify the name or address of the NMS. Specify traps (the default) to send SNMP traps to the host. Specify informs to send SNMP informs to the host. Specify the SNMP version to support. Version 1, the default, is not available with informs. For <i>community-string</i>, specify the string to send with the notification operation. Though you can set this string by using the snmp-server host command, we recommend that you define this string by using the snmp-server community command before using the snmp-server host command. For <i>notification-type</i>, use the mac-notification keyword.
Step 3	snmp-server enable traps mac-notification move	Enable the switch to send MAC address move notification traps to the NMS.
Step 4	mac address-table notification mac-move	Enable the MAC address move notification feature.
Step 5	end	Return to privileged EXEC mode.
Step 6	show mac address-table notification mac-move show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable MAC address-move notification traps, use the **no snmp-server enable traps mac-notification move** global configuration command. To disable the MAC address-move notification feature, use the **no mac address-table notification mac-move** global configuration command.

This example shows how to specify 172.20.10.10 as the NMS, enable the switch to send MAC address move notification traps to the NMS, enable the MAC address move notification feature, and enable traps when a MAC address moves from one port to another.

```
Switch(config)# snmp-server host 172.20.10.10 traps private mac-notification
Switch(config)# snmp-server enable traps mac-notification move
Switch(config)# mac address-table notification mac-move
```

You can verify your settings by entering the **show mac address-table notification mac-move** privileged EXEC commands.

Configuring MAC Threshold Notification Traps

When you configure MAC threshold notification, an SNMP notification is generated and sent to the network management system when a MAC address table threshold limit is reached or exceeded.

Beginning in privileged EXEC mode, follow these steps to configure the switch to send MAC address table threshold 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 notification-type</i>	Specify the recipient of the trap message. <ul style="list-style-type: none"> For <i>host-addr</i>, specify the name or address of the NMS. Specify traps (the default) to send SNMP traps to the host. Specify informs to send SNMP informs to the host. Specify the SNMP version to support. Version 1, the default, is not available with informs. For <i>community-string</i>, specify the string to send with the notification operation. Though you can set this string by using the snmp-server host command, we recommend that you define this string by using the snmp-server community command before using the snmp-server host command. For <i>notification-type</i>, use the mac-notification keyword.
Step 3	snmp-server enable traps mac-notification threshold	Enable the switch to send MAC threshold notification traps to the NMS.
Step 4	mac address-table notification threshold	Enable the MAC address threshold notification feature.
Step 5	mac address-table notification threshold [limit <i>percentage</i>] [interval <i>time</i>]	Enter the threshold value for the MAC address threshold usage monitoring. <ul style="list-style-type: none"> (Optional) For limit percentage, specify the percentage of the MAC address table use; valid values are from 1 to 100 percent. The default is 50 percent. (Optional) For interval time, specify the time between notifications; valid values are greater than or equal to 120 seconds. The default is 120 seconds.
Step 6	end	Return to privileged EXEC mode.
Step 7	show mac address-table notification threshold show running-config	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable MAC address-threshold notification traps, use the **no snmp-server enable traps mac-notification threshold** global configuration command. To disable the MAC address-threshold notification feature, use the **no mac address-table notification threshold** global configuration command.

This example shows how to specify 172.20.10.10 as the NMS, enable the MAC address threshold notification feature, set the interval time to 123 seconds, and set the limit to 78 per cent.

```
Switch(config)# snmp-server host 172.20.10.10 traps private mac-notification
Switch(config)# snmp-server enable traps mac-notification threshold
Switch(config)# mac address-table notification threshold
Switch(config)# mac address-table notification threshold interval 123
Switch(config)# mac address-table notification threshold limit 78
```

You can verify your settings by entering the **show mac address-table notification threshold** privileged EXEC commands.

Adding and Removing Static Address Entries

A static address has these characteristics:

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

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

A 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. <ul style="list-style-type: none"> • 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. Valid VLAN IDs are 1 to 4094. • For <i>interface-id</i>, specify the interface to which the received packet is forwarded. Valid interfaces include physical ports or port channels. For static multicast addresses, you can enter multiple interface IDs. For static unicast addresses, you can enter only one interface at a time, but you can enter the command multiple times with the same MAC address and VLAN ID.

	Command	Purpose
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 port:

```
Switch(config)# mac address-table static c2f3.220a.12f4 vlan 4 interface
gigabitethernet0/1
```

Configuring Unicast MAC Address Filtering

When unicast MAC address filtering is enabled, the switch drops packets with specific source or destination MAC addresses. This feature is disabled by default and only supports unicast static addresses.

Follow these guidelines when using this feature:

- Multicast MAC addresses, broadcast MAC addresses, and router MAC addresses are not supported. If you specify one of these addresses when entering the **mac address-table static *mac-addr* vlan *vlan-id* drop** global configuration command, one of these messages appears:


```
% Only unicast addresses can be configured to be dropped
% CPU destined address cannot be configured as drop address
```
- Packets that are forwarded to the CPU are also not supported.
- If you add a unicast MAC address as a static address and configure unicast MAC address filtering, the switch either adds the MAC address as a static address or drops packets with that MAC address, depending on which command was entered last. The second command that you entered overrides the first command.

For example, if you enter the **mac address-table static *mac-addr* vlan *vlan-id* interface *interface-id*** global configuration command followed by the **mac address-table static *mac-addr* vlan *vlan-id* drop** command, the switch drops packets with the specified MAC address as a source or destination.

If you enter the **mac address-table static *mac-addr* vlan *vlan-id* drop** global configuration command followed by the **mac address-table static *mac-addr* vlan *vlan-id* interface *interface-id*** command, the switch adds the MAC address as a static address.

You enable unicast MAC address filtering and configure the switch to drop packets with a specific address by specifying the source or destination unicast MAC address and the VLAN from which it is received.

Beginning in privileged EXEC mode, follow these steps to configure the switch to drop a source or destination unicast 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> drop	Enable unicast MAC address filtering and configure the switch to drop a packet with the specified source or destination unicast static address. <ul style="list-style-type: none"> For <i>mac-addr</i>, specify a source or destination unicast MAC address. Packets with this MAC address are dropped. For <i>vlan-id</i>, specify the VLAN for which the packet with the specified MAC address is received. Valid VLAN IDs are 1 to 4094.
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 disable unicast MAC address filtering, use the **no mac address-table static *mac-addr* vlan *vlan-id*** global configuration command.

This example shows how to enable unicast MAC address filtering and to configure the switch to drop packets that have a source or destination address of c2f3.220a.12f4. When a packet is received in VLAN 4 with this MAC address as its source or destination, the packet is dropped:

```
Switch(config)# mac address-table static c2f3.220a.12f4 vlan 4 drop
```

Disabling MAC Address Learning on a VLAN

By default, MAC address learning is enabled on all VLANs on the switch. You can control MAC address learning on a VLAN to manage the available MAC address table space by controlling which VLANs, and therefore which ports, can learn MAC addresses. Before you disable MAC address learning, be sure that you are familiar with the network topology and the switch system configuration. Disabling MAC address learning on a VLAN could cause flooding in the network.

Follow these guidelines when disabling MAC address learning on a VLAN:

- Disabling MAC address learning on a VLAN is supported only if the switch is running the IP Services or LAN base image.
- Use caution before disabling MAC address learning on a VLAN with a configured switch virtual interface (SVI). The switch then floods all IP packets in the Layer 2 domain.
- You can disable MAC address learning on a single VLAN ID (for example, **no mac address-table learning vlan 223**) or on a range of VLAN IDs (for example, **no mac address-table learning vlan 1-20, 15**).
- We recommend that you disable MAC address learning only in VLANs with two ports. If you disable MAC address learning on a VLAN with more than two ports, every packet entering the switch is flooded in that VLAN domain.
- You cannot disable MAC address learning on a VLAN that is used internally by the switch. If the VLAN ID that you enter is an internal VLAN, the switch generates an error message and rejects the command. To view internal VLANs in use, enter the **show vlan internal usage** privileged EXEC command.

- If you disable MAC address learning on a VLAN configured as a private-VLAN primary VLAN, MAC addresses are still learned on the secondary VLAN that belongs to the private VLAN and are then replicated on the primary VLAN. If you disable MAC address learning on the secondary VLAN, but not the primary VLAN of a private VLAN, MAC address learning occurs on the primary VLAN and is replicated on the secondary VLAN.
- You cannot disable MAC address learning on an RSPAN VLAN. The configuration is not allowed.
- If you disable MAC address learning on a VLAN that includes a secure port, MAC address learning is not disabled on that port. If you disable port security, the configured MAC address learning state is enabled.

Beginning in privileged EXEC mode, follow these steps to disable MAC address learning on a VLAN:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no mac address-table learning vlan <i>vlan-id</i>	Disable MAC address learning on the specified VLAN or VLANs. You can specify a single VLAN ID or a range of VLAN IDs separated by a hyphen or comma. Valid VLAN IDs are 1 to 4094.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mac address-table learning [vlan <i>vlan-id]</i>	Verify the configuration.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To reenabling MAC address learning on a VLAN, use the **default mac address-table learning vlan** *vlan-id* global configuration command. You can also reenabling MAC address learning on a VLAN by entering the **mac address-table learning vlan** *vlan-id* global configuration command. The first (**default**) command returns to a default condition and therefore does not appear in the output from the **show running-config** command. The second command causes the configuration to appear in the **show running-config** privileged EXEC command display.

This example shows how to disable MAC address learning on VLAN 200:

```
Switch(config)# no mac address-table learning vlan 200
```

You can display the MAC address learning status of all VLANs or a specified VLAN by entering the **show mac-address-table learning [vlan** *vlan-id]* privileged EXEC command.

Displaying Address Table Entries

You can display the MAC address table by using one or more of the privileged EXEC commands described in [Table 1-3](#):

Table 1-3 Commands for Displaying the MAC Address Table

Command	Description
show ip igmp snooping groups	Displays the Layer 2 multicast entries for all VLANs or the specified VLAN.
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.

Table 1-3 Commands for Displaying the MAC Address Table (continued)

Command	Description
<code>show mac address-table dynamic</code>	Displays only dynamic MAC address table entries.
<code>show mac address-table interface</code>	Displays the MAC address table information for the specified interface.
<code>show mac address-table learning</code>	Displays MAC address learning status of all VLANs or the specified VLAN.
<code>show mac address-table notification</code>	Displays the MAC notification parameters and history table.
<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 learn the 48-bit MAC address or the local data link address of that device. The process of learning 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. Using an IP address, ARP finds the associated MAC address. When a MAC address is found, 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.



Note

For CLI procedures, see the Cisco IOS Release 12.4 documentation on Cisco.com.



CHAPTER 1

Configuring Web-Based Authentication

This chapter describes how to configure web-based authentication. It contains these sections:

- [Understanding Web-Based Authentication, page 1-1](#)
- [Configuring Web-Based Authentication, page 1-9](#)
- [Displaying Web-Based Authentication Status, page 1-17](#)



Note

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

Understanding Web-Based Authentication

Use the web-based authentication feature, known as *web authentication proxy*, to authenticate end users on host systems that do not run the IEEE 802.1x supplicant.



Note

You can configure web-based authentication on Layer 2 and Layer 3 interfaces.

When you initiate an HTTP session, web-based authentication intercepts ingress HTTP packets from the host and sends an HTML login page to the users. The users enter their credentials, which the web-based authentication feature sends to the authentication, authorization, and accounting (AAA) server for authentication.

If authentication succeeds, web-based authentication sends a Login-Successful HTML page to the host and applies the access policies returned by the AAA server.

If authentication fails, web-based authentication forwards a Login-Fail HTML page to the user, prompting the user to retry the login. If the user exceeds the maximum number of attempts, web-based authentication forwards a Login-Expired HTML page to the host, and the user is placed on a watch list for a waiting period.

These sections describe the role of web-based authentication as part of AAA:

- [Device Roles, page 1-2](#)
- [Host Detection, page 1-2](#)
- [Session Creation, page 1-3](#)
- [Authentication Process, page 1-3](#)

- [Web Authentication Customizable Web Pages, page 1-6](#)
- [Web-based Authentication Interactions with Other Features, page 1-7](#)

Device Roles

With web-based authentication, the devices in the network have these specific roles:

- *Client*—The device (workstation) that requests access to the LAN and the services and responds to requests from the switch. The workstation must be running an HTML browser with Java Script enabled.
- *Authentication server*—Authenticates the client. The authentication server validates the identity of the client and notifies the switch that the client is authorized to access the LAN and the switch services or that the client is denied.
- *Switch*—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.

Figure 1-1 shows the roles of these devices in a network:

Figure 1-1 Web-Based Authentication Device Roles

Host Detection

The switch maintains an IP device tracking table to store information about detected hosts.



Note

By default, the IP device tracking feature is disabled on a switch. You must enable the IP device tracking feature to use web-based authentication.

For Layer 2 interfaces, web-based authentication detects IP hosts by using these mechanisms:

- ARP based trigger—ARP redirect ACL allows web-based authentication to detect hosts with a static IP address or a dynamic IP address.
- Dynamic ARP inspection
- DHCP snooping—Web-based authentication is notified when the switch creates a DHCP-binding entry for the host.

Session Creation

When web-based authentication detects a new host, it creates a session as follows:

- Reviews the exception list.
If the host IP is included in the exception list, the policy from the exception list entry is applied, and the session is established.
- Reviews for authorization bypass
If the host IP is not on the exception list, web-based authentication sends a nonresponsive-host (NRH) request to the server.
If the server response is *access accepted*, authorization is bypassed for this host. The session is established.
- Sets up the HTTP intercept ACL
If the server response to the NRH request is *access rejected*, the HTTP intercept ACL is activated, and the session waits for HTTP traffic from the host.

Authentication Process

When you enable web-based authentication, these events occur:

- The user initiates an HTTP session.
- The HTTP traffic is intercepted, and authorization is initiated. The switch sends the login page to the user. The user enters a username and password, and the switch sends the entries to the authentication server.
- If the authentication succeeds, the switch downloads and activates the user's access policy from the authentication server. The login success page is sent to the user.
- If the authentication fails, the switch sends the login fail page. The user retries the login. If the maximum number of attempts fails, the switch sends the login expired page, and the host is placed in a watch list. After the watch list times out, the user can retry the authentication process.
- If the authentication server does not respond to the switch, and if an AAA fail policy is configured, the switch applies the failure access policy to the host. The login success page is sent to the user. (See the [“Local Web Authentication Banner”](#) section on page 1-4.)
- The switch reauthenticates a client when the host does not respond to an ARP probe on a Layer 2 interface, or when the host does not send any traffic within the idle timeout on a Layer 3 interface.
- The feature applies the downloaded timeout or the locally configured session timeout.
- If the terminate action is RADIUS, the feature sends a nonresponsive host (NRH) request to the server. The terminate action is included in the response from the server.
- If the terminate action is default, the session is dismantled, and the applied policy is removed.

Local Web Authentication Banner


You can create a banner that will appear when you log in to a switch by using web authentication.

The banner appears on both the login page and the authentication-result pop-up pages.

- *Authentication Successful*
- *Authentication Failed*
- *Authentication Expired*

You create a banner by using the **ip admission auth-proxy-banner http** global configuration command. The default banner *Cisco Systems* and *Switch host-name Authentication* appear on the Login Page. *Cisco Systems* appears on the authentication result pop-up page, as shown in [Figure 1-2](#).


Figure 1-2 Authentication Successful Banner



You can also customize the banner, as shown in [Figure 1-3](#).

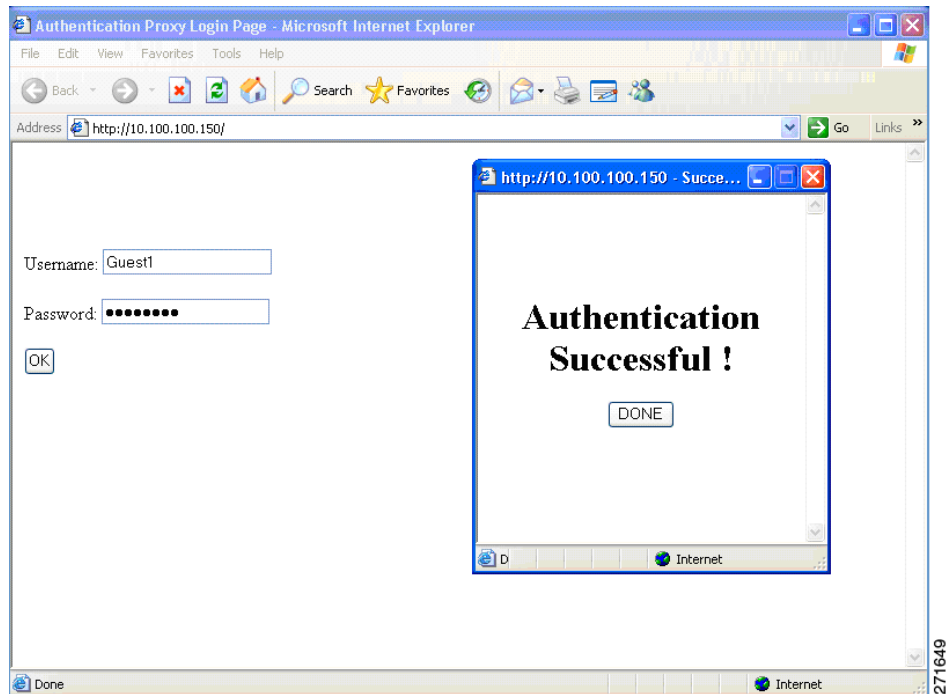
- Add a switch, router, or company name to the banner by using the **ip admission auth-proxy-banner http banner-text** global configuration command.
- Add a logo or text file to the banner by using the **ip admission auth-proxy-banner http file-path** global configuration command.

Figure 1-3 Customized Web Banner



If you do not enable a banner, only the username and password dialog boxes appear in the web authentication login screen, and no banner appears when you log into the switch, as shown in Figure 1-4.

Figure 1-4 Login Screen With No Banner



For more information, see the [Cisco IOS Security Command Reference](#) and the "Configuring a Web Authentication Local Banner" section on page 1-16.

Web Authentication Customizable Web Pages

During the web-based authentication process, the switch internal HTTP server hosts four HTML pages to deliver to an authenticating client. The server uses these pages to notify you of these four-authentication process states:


- Login—Your credentials are requested.
- Success—The login was successful.
- Fail—The login failed.
- Expire—The login session has expired because of excessive login failures.

Guidelines

- You can substitute your own HTML pages for the default internal HTML pages.
- You can use a logo or specify text in the *login*, *success*, *failure*, and *expire* web pages.
- On the banner page, you can specify text in the login page.
- The pages are in HTML.
- You must include an HTML redirect command in the success page to access a specific URL.
- The URL string must be a valid URL (for example, <http://www.cisco.com>). An incomplete URL might cause *page not found* or similar errors on a web browser.
- If you configure web pages for HTTP authentication, they must include the appropriate HTML commands (for example, to set the page time out, to set a hidden password, or to confirm that the same page is not submitted twice).
- The CLI command to redirect users to a specific URL is not available when the configured login form is enabled. The administrator should ensure that the redirection is configured in the web page.
- If the CLI command redirecting users to specific URL after authentication occurs is entered and then the command configuring web pages is entered, the CLI command redirecting users to a specific URL does not take effect.
- Configured web pages can be copied to the switch boot flash or flash.
- Configured pages can be accessed from the flash on the stack master or members.
- The login page can be on one flash, and the success and failure pages can be another flash (for example, the flash on the stack master or a member).
- You must configure all four pages.
- The banner page has no effect if it is configured with the web page.
- All of the logo files (image, flash, audio, video, and so on) that are stored in the system directory (for example, flash, disk0, or disk) and that must be displayed on the login page must use `web_auth_<filename>` as the file name.
- The configured authentication proxy feature supports both HTTP and SSL.

You can substitute your HTML pages, as shown in [Figure 1-5 on page 1-7](#), for the default internal HTML pages. You can also specify a URL to which users are redirected after authentication occurs, which replaces the internal Success page.

Figure 1-5 Customizable Authentication Page



For more information, see the [“Customizing the Authentication Proxy Web Pages”](#) section on page 1-14.

Web-based Authentication Interactions with Other Features

- [Port Security, page 1-7](#)
- [LAN Port IP, page 1-7](#)
- [Gateway IP, page 1-8](#)
- [ACLs, page 1-8](#)
- [Context-Based Access Control, page 1-8](#)
- [802.1x Authentication, page 1-8](#)
- [EtherChannel, page 1-8](#)

Port Security

You can configure web-based authentication and port security on the same port. Web-based authentication 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 the port.

For more information about enabling port security, see the [“Configuring Port Security”](#) section on page 1-8.

LAN Port IP

You can configure LAN port IP (LPI) and Layer 2 web-based authentication on the same port. The host is authenticated by using web-based authentication first, followed by LPI posture validation. The LPI host policy overrides the web-based authentication host policy.

If the web-based authentication idle timer expires, the NAC policy is removed. The host is authenticated, and posture is validated again.

Gateway IP

You cannot configure Gateway IP (GWIP) on a Layer 3 VLAN interface if web-based authentication is configured on any of the switch ports in the VLAN.

You can configure web-based authentication on the same Layer 3 interface as Gateway IP. The host policies for both features are applied in software. The GWIP policy overrides the web-based authentication host policy.

ACLs

If you configure a VLAN ACL or a Cisco IOS ACL on an interface, the ACL is applied to the host traffic only after the web-based authentication host policy is applied.

For Layer 2 web-based authentication, you must configure a port ACL (PACL) as the default access policy for ingress traffic from hosts connected to the port. After authentication, the web-based authentication host policy overrides the PACL.



Note

When a proxy ACL is configured for a web-based authentication client, the proxy ACL is downloaded and applied as part of the authorization process. Hence, the PACL displays the proxy ACL access control entry (ACE).

You cannot configure a MAC ACL and web-based authentication on the same interface.

You cannot configure web-based authentication on a port whose access VLAN is configured for VACL capture.

Context-Based Access Control

Web-based authentication cannot be configured on a Layer 2 port if context-based access control (CBAC) is configured on the Layer 3 VLAN interface of the port VLAN.

802.1x Authentication

You cannot configure web-based authentication on the same port as 802.1x authentication except as a fallback authentication method.

EtherChannel

You can configure web-based authentication on a Layer 2 EtherChannel interface. The web-based authentication configuration applies to all member channels.

Configuring Web-Based Authentication

- [Default Web-Based Authentication Configuration, page 1-9](#)
- [Web-Based Authentication Configuration Guidelines and Restrictions, page 1-9](#)
- [Web-Based Authentication Configuration Task List, page 1-10](#)
- [Configuring the Authentication Rule and Interfaces, page 1-10](#)
- [Configuring AAA Authentication, page 1-11](#)
- [Configuring Switch-to-RADIUS-Server Communication, page 1-12](#)
- [Configuring the HTTP Server, page 1-13](#)
- [Configuring the Web-Based Authentication Parameters, page 1-16](#)
- [Removing Web-Based Authentication Cache Entries, page 1-17](#)

Default Web-Based Authentication Configuration

Table 1-1 shows the default web-based authentication configuration.

Table 1-1 Default Web-based Authentication Configuration

Feature	Default Setting
AAA	Disabled
RADIUS server	<ul style="list-style-type: none"> • None specified
<ul style="list-style-type: none"> • IP address • UDP authentication port • Key 	<ul style="list-style-type: none"> • 1812 • None specified
Default value of inactivity timeout	3600 seconds
Inactivity timeout	Enabled

Web-Based Authentication Configuration Guidelines and Restrictions

- Web-based authentication is an ingress-only feature.
- You can configure web-based authentication only on access ports. Web-based authentication is not supported on trunk ports, EtherChannel member ports, or dynamic trunk ports.
- You must configure the default ACL on the interface before configuring web-based authentication. Configure a port ACL for a Layer 2 interface or a Cisco IOS ACL for a Layer 3 interface.
- You cannot authenticate hosts on Layer 2 interfaces with static ARP cache assignment. These hosts are not detected by the web-based authentication feature because they do not send ARP messages.
- By default, the IP device tracking feature is disabled on a switch. You must enable the IP device tracking feature to use web-based authentication.
- You must configure at least one IP address to run the switch HTTP server. You must also configure routes to reach each host IP address. The HTTP server sends the HTTP login page to the host.

- Hosts that are more than one hop away might experience traffic disruption if an STP topology change results in the host traffic arriving on a different port. This occurs because the ARP and DHCP updates might not be sent after a Layer 2 (STP) topology change.
- Web-based authentication does not support VLAN assignment as a downloadable-host policy.
- Web-based authentication is not supported for IPv6 traffic.
- Web-based authentication and Network Edge Access Topology (NEAT) are mutually exclusive. You cannot use web-based authentication when NEAT is enabled on an interface, and you cannot use NEAT when web-based authentication is running on an interface.
- Web-based authentication supports only RADIUS authorization servers. You cannot use TACACS+ servers or local authorization.

Web-Based Authentication Configuration Task List

- [Configuring the Authentication Rule and Interfaces, page 1-10](#)
- [Configuring AAA Authentication, page 1-11](#)
- [Configuring Switch-to-RADIUS-Server Communication, page 1-12](#)
- [Configuring the HTTP Server, page 1-13](#)
- [Configuring the Web-Based Authentication Parameters, page 1-16](#)
- [Configuring the Web-Based Authentication Parameters, page 1-16](#)
- [Removing Web-Based Authentication Cache Entries, page 1-17](#)

Configuring the Authentication Rule and Interfaces

	Command	Purpose
Step 1	ip admission name <i>name</i> proxy http	Configure an authentication rule for web-based authorization.
Step 2	interface <i>type slot/port</i>	Enter interface configuration mode and specifies the ingress Layer 2 or Layer 3 interface to be enabled for web-based authentication. <i>type</i> can be fastethernet, gigabit ethernet, or tengigabitethernet.
Step 3	ip access-group <i>name</i>	Apply the default ACL.
Step 4	ip admission <i>name</i>	Configures web-based authentication on the specified interface.
Step 5	exit	Return to configuration mode.
Step 6	ip device tracking	Enables the IP device tracking table.
Step 7	end	Return to privileged EXEC mode.
Step 8	show ip admission configuration	Display the configuration.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to enable web-based authentication on Fast Ethernet port 5/1:

```
Switch(config)# ip admission name webauth1 proxy http
Switch(config)# interface fastethernet 5/1
Switch(config-if)# ip admission webauth1
Switch(config-if)# exit
Switch(config)# ip device tracking
```

This example shows how to verify the configuration:

```
Switch# show ip admission configuration
Authentication Proxy Banner not configured
Authentication global cache time is 60 minutes
Authentication global absolute time is 0 minutes
Authentication global init state time is 2 minutes
Authentication Proxy Watch-list is disabled

Authentication Proxy Rule Configuration
Auth-proxy name webauth1
http list not specified inactivity-time 60 minutes

Authentication Proxy Auditing is disabled
Max Login attempts per user is 5
```

Configuring AAA Authentication

	Command	Purpose
Step 1	aaa new-model	Enables AAA functionality.
Step 2	aaa authentication login default group { <i>tacacs+</i> <i>radius</i> }	Defines the list of authentication methods at login.
Step 3	aaa authorization auth-proxy default group { <i>tacacs+</i> <i>radius</i> }	Create an authorization method list for web-based authorization.
Step 4	radius-server host { <i>hostname</i> <i>ip-address</i> } test username <i>username</i>	Specify an AAA server. Specify the host name or IP address of the remote RADIUS server. The test username <i>username</i> option enables automated testing of the RADIUS server connection. The specified username does not need to be a valid user name.
Step 5	radius-server key <i>string</i>	Configure the authorization and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. To use multiple RADIUS servers, reenter this command for each server.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to enable AAA:

```
Switch(config)# aaa new-model
Switch(config)# aaa authentication login default group radius
Switch(config)# aaa authorization auth-proxy default group radius
```

Configuring Switch-to-RADIUS-Server Communication

RADIUS security servers identification:

- Host name
- Host IP address
- Host name and specific UDP port numbers
- IP address and specific UDP port numbers

The combination of the IP address and UDP port number creates a unique identifier, that 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 that is configured functions as the failover backup to the first one. The RADIUS host entries are chosen in the order that they were configured.

To configure the RADIUS server parameters, perform this task:

	Command	Purpose
Step 1	ip radius source-interface <i>interface_name</i>	Specify that the RADIUS packets have the IP address of the indicated interface.
Step 2	radius-server host { <i>hostname</i> <i>ip-address</i> } test username <i>username</i>	Specify the host name or IP address of the remote RADIUS server. The test username <i>username</i> option enables automated testing of the RADIUS server connection. The specified <i>username</i> does not need to be a valid user name. The key option specifies an authentication and encryption key to use between the switch and the RADIUS server. To use multiple RADIUS servers, reenter this command for each server.
Step 3	radius-server key <i>string</i>	Configure the authorization and encryption key used between the switch and the RADIUS daemon running on the RADIUS server.
Step 4	radius-server vsa send authentication	Enable downloading of an ACL from the RADIUS server. This feature is supported in Cisco IOS Release 12.2(50)SG.
Step 5	radius-server dead-criteria tries <i>num-tries</i>	Specify the number of unanswered sent messages to a RADIUS server before considering the server to be inactive. The range of <i>num-tries</i> is 1 to 100.

When you configure the RADIUS server parameters:

- Specify the **key** *string* on a separate command line.
- For **key** *string*, 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.
- When you specify the **key** *string*, use spaces within and at the end of the key. 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.

- You can globally configure the timeout, retransmission, and encryption key values for all RADIUS servers by using with 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 *Cisco IOS Security Configuration Guide* and the *Cisco IOS Security Command Reference*: http://www.cisco.com/en/US/docs/ios/12_2/security/command/reference/fsecur_r.html

**Note**

You need to configure some settings on the RADIUS server, including: the switch IP address, the key string to be shared by both the server and the switch, and the downloadable ACL (DACL). For more information, see the RADIUS server documentation.

This example shows how to configure the RADIUS server parameters on a switch:

```
Switch(config)# ip radius source-interface Vlan80
Switch(config)# radius-server host 172.120.39.46 test username user1
Switch(config)# radius-server key rad123
Switch(config)# radius-server dead-criteria tries 2
```

Configuring the HTTP Server

To use web-based authentication, you must enable the HTTP server within the switch. You can enable the server for either HTTP or HTTPS.

	Command	Purpose
Step 1	ip http server	Enable the HTTP server. The web-based authentication feature uses the HTTP server to communicate with the hosts for user authentication.
Step 2	ip http secure-server	Enable HTTPS.

You can configure custom authentication proxy web pages or specify a redirection URL for successful login.

**Note**

To ensure secure authentication when you enter the **ip http secure-server** command, the login page is always in HTTPS (secure HTTP) even if the user sends an HTTP request.

- [Customizing the Authentication Proxy Web Pages](#)
- [Specifying a Redirection URL for Successful Login](#)

Customizing the Authentication Proxy Web Pages

You can configure web authentication to display four substitute HTML pages to the user in place of the switch default HTML pages during web-based authentication.

To specify the use of your custom authentication proxy web pages, first store your custom HTML files on the switch flash memory, then perform this task in global configuration mode:

	Command	Purpose
Step 1	ip admission proxy http login page file <i>device:login-filename</i>	Specify the location in the switch memory file system of the custom HTML file to use in place of the default login page. The <i>device:</i> is flash memory.
Step 2	ip admission proxy http success page file <i>device:success-filename</i>	Specify the location of the custom HTML file to use in place of the default login success page.
Step 3	ip admission proxy http failure page file <i>device:fail-filename</i>	Specify the location of the custom HTML file to use in place of the default login failure page.
Step 4	ip admission proxy http login expired page file <i>device:expired-filename</i>	Specify the location of the custom HTML file to use in place of the default login expired page.

When configuring customized authentication proxy web pages, follow these guidelines:

- To enable the custom web pages feature, specify all four custom HTML files. If you specify fewer than four files, the internal default HTML pages are used.
- The four custom HTML files must be present on the flash memory of the switch. The maximum size of each HTML file is 8 KB.
- Any images on the custom pages must be on an accessible HTTP server. Configure an intercept ACL within the admission rule.
- Any external link from a custom page requires configuration of an intercept ACL within the admission rule.
- To access a valid DNS server, any name resolution required for external links or images requires configuration of an intercept ACL within the admission rule.
- If the custom web pages feature is enabled, a configured auth-proxy-banner is not used.
- If the custom web pages feature is enabled, the redirection URL for successful login feature is not available.
- To remove the specification of a custom file, use the **no** form of the command.

Because the custom login page is a public web form, consider these guidelines for the page:

- The login form must accept user entries for the username and password and must show them as **uname** and **pwd**.
- The custom login page should follow best practices for a web form, such as page timeout, hidden password, and prevention of redundant submissions.

This example shows how to configure custom authentication proxy web pages:

```
Switch(config)# ip admission proxy http login page file flash:login.htm
Switch(config)# ip admission proxy http success page file flash:success.htm
Switch(config)# ip admission proxy http fail page file flash:fail.htm
Switch(config)# ip admission proxy http login expired page flash:expired.htm
```

This example shows how to verify the configuration of a custom authentication proxy web pages:

```
Switch# show ip admission configuration
Authentication proxy webpage
  Login page           : flash:login.htm
  Success page        : flash:success.htm
  Fail Page           : flash:fail.htm
  Login expired Page  : flash:expired.htm

Authentication global cache time is 60 minutes
Authentication global absolute time is 0 minutes
Authentication global init state time is 2 minutes
Authentication Proxy Session ratelimit is 100
Authentication Proxy Watch-list is disabled
Authentication Proxy Auditing is disabled
Max Login attempts per user is 5
```

Specifying a Redirection URL for Successful Login

You can specify a URL to which the user is redirected after authentication, effectively replacing the internal *Success* HTML page.

Command	Purpose
<code>ip admission proxy http success redirect <i>url-string</i></code>	Specify a URL for redirection of the user in place of the default login success page.

When configuring a redirection URL for successful login, consider these guidelines:

- If the custom authentication proxy web pages feature is enabled, the redirection URL feature is disabled and is not available in the CLI. You can perform redirection in the custom-login success page.
- If the redirection URL feature is enabled, a configured auth-proxy-banner is not used.
- To remove the specification of a redirection URL, use the **no** form of the command.

This example shows how to configure a redirection URL for successful login:

```
Switch(config)# ip admission proxy http success redirect www.cisco.com
```

This example shows how to verify the redirection URL for successful login:

```
Switch# show ip admission configuration
Authentication Proxy Banner not configured
Customizable Authentication Proxy webpage not configured
HTTP Authentication success redirect to URL: http://www.cisco.com
Authentication global cache time is 60 minutes
Authentication global absolute time is 0 minutes
Authentication global init state time is 2 minutes
Authentication Proxy Watch-list is disabled
Authentication Proxy Max HTTP process is 7
Authentication Proxy Auditing is disabled
Max Login attempts per user is 5
```

Configuring the Web-Based Authentication Parameters

You can configure the maximum number of failed login attempts before the client is placed in a watch list for a waiting period.

	Command	Purpose
Step 1	ip admission max-login-attempts <i>number</i>	Set the maximum number of failed login attempts. The range is 1 to 2147483647 attempts. The default is 5.
Step 2	end	Returns to privileged EXEC mode.
Step 3	show ip admission configuration	Display the authentication proxy configuration.
Step 4	show ip admission cache	Display the list of authentication entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to set the maximum number of failed login attempts to 10:

```
Switch(config)# ip admission max-login-attempts 10
```

Configuring a Web Authentication Local Banner

Beginning in privileged EXEC mode, follow these steps to configure a local banner on a switch that has web authentication configured.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip admission auth-proxy-banner http <i>[banner-text file-path]</i>	Enable the local banner. (Optional) Create a custom banner by entering <i>C banner-text C</i> , where <i>C</i> is a delimiting character or a file-path indicates a file (for example, a logo or text file) that appears in the banner.
Step 3	end	Return to privileged EXEC mode.
Step 4	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to configure a local banner with the custom message *My Switch*:

```
Switch(config) configure terminal
Switch(config)# aaa new-model
Switch(config)# aaa ip auth-proxy auth-proxy-banner C My Switch C
Switch(config) end
```

For more information about the **ip auth-proxy auth-proxy-banner** command, see the “Authentication Proxy Commands” section of the [Cisco IOS Security Command Reference](#) on Cisco.com.

Removing Web-Based Authentication Cache Entries

Command	Purpose
<code>clear ip auth-proxy cache { * host ip address }</code>	Delete authentication proxy entries. Use an asterisk to delete all cache entries. Enter a specific IP address to delete the entry for a single host.
<code>clear ip admission cache { * host ip address }</code>	Delete authentication proxy entries. Use an asterisk to delete all cache entries. Enter a specific IP address to delete the entry for a single host.

This example shows how to remove the web-based authentication session for the client at the IP address 209.165.201.1:

```
Switch# clear ip auth-proxy cache 209.165.201.1
```

Displaying Web-Based Authentication Status

Perform this task to display the web-based authentication settings for all interfaces or for specific ports:

Command	Purpose
Step 1 <code>show authentication sessions</code> <code>[interface type slot/port]</code>	Displays the web-based authentication settings. type = fastethernet, gigabitethernet, or tengigabitethernet (Optional) Use the interface keyword to display the web-based authentication settings for a specific interface.

This example shows how to view only the global web-based authentication status:

```
Switch# show authentication sessions
```

This example shows how to view the web-based authentication settings for gigabit interface 3/27:

```
Switch# show authentication sessions interface gigabitethernet 3/27
```




CHAPTER 1

Clustering Switches

This chapter provides the concepts and procedures to create and manage Catalyst 2960 and 2960-S switch clusters. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

You can create and manage switch clusters by using Cisco Network Assistant (hereafter known as Network Assistant), the command-line interface (CLI), or SNMP. For complete procedures, see the online help. For the CLI cluster commands, see the switch command reference.



Note

Network Assistant supports switch clusters, but we recommend that you instead group switches into *communities*. Network Assistant has a Cluster Conversion Wizard to help you convert a cluster to a *community*. For more information about Network Assistant, including introductory information on managing switch clusters and converting a switch cluster to a community, see *Getting Started with Cisco Network Assistant*, available on Cisco.com.

This chapter focuses on Catalyst 2960 and 2960-S switch clusters. It also includes guidelines and limitations for clusters mixed with other cluster-capable Catalyst switches, but it does not provide complete descriptions of the cluster features for these other switches. For complete cluster information for a specific Catalyst platform, refer to the software configuration guide for that switch.

This chapter consists of these sections:

- [Understanding Switch Clusters, page 1-2](#)
- [Planning a Switch Cluster, page 1-4](#)
- [Using the CLI to Manage Switch Clusters, page 1-15](#)
- [Using SNMP to Manage Switch Clusters, page 1-16](#)



Note

We do not recommend using the **ip http access-class** global configuration command to limit access to specific hosts or networks. Access should be controlled through the cluster command switch or by applying access control lists (ACLs) on interfaces that are configured with IP address. For more information on ACLs, see [Chapter 1, “Configuring Network Security with ACLs.”](#)

Understanding Switch Clusters

A *switch cluster* is a set of up to 16 connected, cluster-capable Catalyst switches that are managed as a single entity. The switches in the cluster use the switch clustering technology so that you can configure and troubleshoot a group of different Catalyst desktop switch platforms through a single IP address.

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



Note

A switch cluster is different from a *switch stack*. A switch stack is a set of Catalyst 2960-S switches connected through their stack ports. For more information about how switch stacks differ from switch clusters, see the “[Switch Clusters and Switch Stacks](#)” section on page 1-13. Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

The benefits of clustering switches include:

- Management of Catalyst switches regardless of their interconnection media and their physical locations. The switches can be in the same location, or they can be distributed across a Layer 2 or Layer 3 (if your cluster is using a Catalyst 3550, Catalyst 3560, or Catalyst 3750 switch as a Layer 3 router between the Layer 2 switches in the cluster) network.

Cluster members are connected to the cluster command switch according to the connectivity guidelines described in the “[Automatic Discovery of Cluster Candidates and Members](#)” section on page 1-5. This section includes management VLAN considerations for the Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL switches. For complete information about these switches in a switch-cluster environment, refer to the software configuration guide for that specific switch.

- Command-switch redundancy if a cluster command switch fails. One or more switches can be designated as *standby cluster command switches* to avoid loss of contact with cluster members. A *cluster standby group* is a group of standby cluster command switches.
- Management of a variety of Catalyst switches through a single IP address. This conserves on IP addresses, especially if you have a limited number of them. All communication with the switch cluster is through the cluster command switch IP address.

[Table 1-1](#) lists the Catalyst switches eligible for switch clustering, including which ones can be cluster command switches and which ones can only be cluster member switches, and the required software versions.

Table 1-1 Switch Software and Cluster Capability

Switch	Cisco IOS Release	Cluster Capability
Catalyst 3750-X or Catalyst 3560-X	12.2(53)SE2 or later	Member or command switch
Catalyst 3750-E or Catalyst 3560-E	12.2(35)SE2 or later	Member or command switch
Catalyst 3750	12.1(11)AX or later	Member or command switch
Catalyst 3560	12.1(19)EA1b or later	Member or command switch
Catalyst 3550	12.1(4)EA1 or later	Member or command switch
Catalyst 2975	12.2(46)EX or later	Member or command switch
Catalyst 2970	12.1(11)AX or later	Member or command switch

Table 1-1 Switch Software and Cluster Capability (continued)

Switch	Cisco IOS Release	Cluster Capability
Catalyst 2960-S	12.2(53)SE or later	Member or command switch
Catalyst 2960	12.2(25)FX or later	Member or command switch
Catalyst 2955	12.1(12c)EA1 or later	Member or command switch
Catalyst 2950	12.0(5.2)WC(1) or later	Member or command switch
Catalyst 2950 LRE	12.1(11)JY or later	Member or command switch
Catalyst 2940	12.1(13)AY or later	Member or command switch
Catalyst 3500 XL	12.0(5.1)XU or later	Member or command switch
Catalyst 2900 XL (8-MB switches)	12.0(5.1)XU or later	Member or command switch
Catalyst 2900 XL (4-MB switches)	11.2(8.5)SA6 (recommended)	Member switch only
Catalyst 1900 and 2820	9.00(-A or -EN) or later	Member switch only

Cluster Command Switch Characteristics

A cluster command switch must meet these requirements:

- It is running Cisco IOS Release 12.2(25)FX or later for a Catalyst 2960 switch, or Cisco IOS Release 12.2(53)SE or later for a Catalyst 2960-S switch.
- It has an IP address.
- It has Cisco Discovery Protocol (CDP) version 2 enabled (the default).
- It is not a command or cluster member switch of another cluster.
- It is connected to the standby cluster command switches through the management VLAN and to the cluster member switches through a common VLAN.

Standby Cluster Command Switch Characteristics

A standby cluster command switch must meet these requirements:

- It is running Cisco IOS 12.2(25)FX or later for a Catalyst 2960 switch, or Cisco IOS Release 12.2(53)SE or later for a Catalyst 2960-S switch.
- It has an IP address.
- It has CDP version 2 enabled.
- It is connected to the command switch and to other standby command switches through its management VLAN.
- It is connected to all other cluster member switches (except the cluster command and standby command switches) through a common VLAN.
- It is redundantly connected to the cluster so that connectivity to cluster member switches is maintained.
- It is not a command or member switch of another cluster.

**Note**

Standby cluster command switches must be the same type of switches as the cluster command switch. For example, if the cluster command switch is a Catalyst 2960 switch, the standby cluster command switches must also be Catalyst 2960 switches. If the cluster command switch is a Catalyst 2960-S switch, the standby cluster command switches must also be Catalyst 2960-S switches. Refer to the switch configuration guide of other cluster-capable switches for their requirements on standby cluster command switches.

Candidate Switch and Cluster Member Switch Characteristics

Candidate switches are cluster-capable switches and switch stacks that have not yet been added to a cluster. Cluster member switches are switches and switch stacks that have actually been added to a switch cluster. Although not required, a candidate or cluster member switch can have its own IP address and password (for related considerations, see the “[IP Addresses](#)” section on page 1-12 and “[Passwords](#)” section on page 1-13).

**Note**

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

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

- It is running cluster-capable software.
- It has CDP version 2 enabled.
- It is not a command or cluster member switch of another cluster.
- If a cluster standby group exists, it is connected to every standby cluster command switch through at least one common VLAN. The VLAN to each standby cluster command switch can be different.
- It is connected to the cluster command switch through at least one common VLAN.

**Note**

Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL candidate and cluster member switches must be connected through their management VLAN to the cluster command switch and standby cluster command switches. For complete information about these switches in a switch-cluster environment, refer to the software configuration guide for that specific switch.

This requirement does not apply if you have a Catalyst 2970, Catalyst 3550, Catalyst 3560, or Catalyst 3750 cluster command switch. Candidate and cluster member switches can connect through any VLAN in common with the cluster command switch.

Planning a Switch Cluster

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

- [Automatic Discovery of Cluster Candidates and Members, page 1-5](#)
- [HSRP and Standby Cluster Command Switches, page 1-9](#)
- [IP Addresses, page 1-12](#)

- [Hostnames, page 1-12](#)
- [Passwords, page 1-13](#)
- [SNMP Community Strings, page 1-13](#)
- [Switch Clusters and Switch Stacks, page 1-13](#)
- [TACACS+ and RADIUS, page 1-15](#)
- [LRE Profiles, page 1-15](#)

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

Automatic Discovery of Cluster Candidates and Members

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



Note

Do not disable CDP on the cluster command switch, on cluster members, or on any cluster-capable switches that you might want a cluster command switch to discover. For more information about CDP, see [Chapter 1, “Configuring CDP.”](#)

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

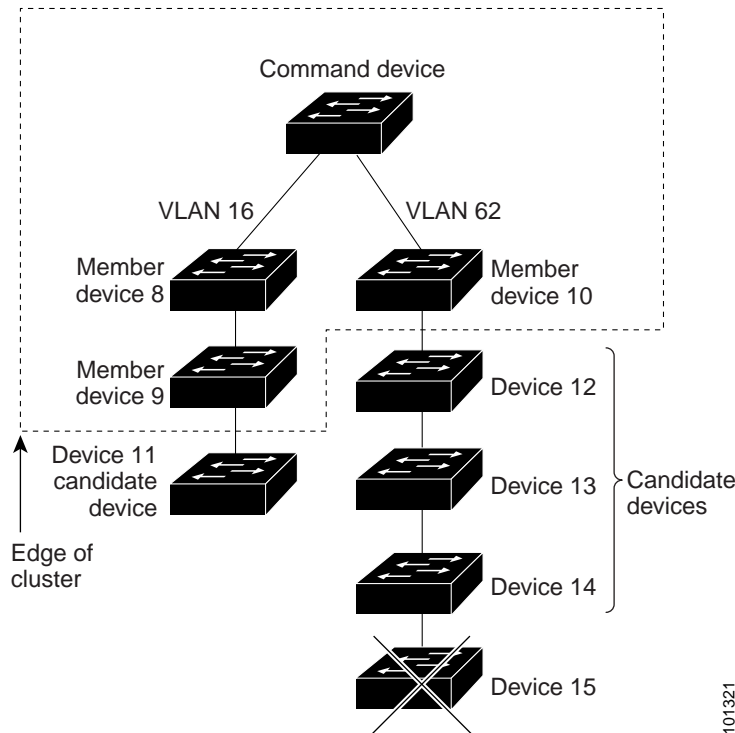
- [Discovery Through CDP Hops, page 1-5](#)
- [Discovery Through Non-CDP-Capable and Noncluster-Capable Devices, page 1-6](#)
- [Discovery Through Different VLANs, page 1-7](#)
- [Discovery Through Different Management VLANs, page 1-7](#)
- [Discovery of Newly Installed Switches, page 1-8](#)

Discovery Through CDP Hops

By using CDP, a cluster command switch can discover switches up to seven CDP hops away (the default is three hops) from the edge of the cluster. The edge of the cluster is where the last cluster member switches are connected to the cluster and to candidate switches. For example, cluster member switches 9 and 10 in [Figure 1-1](#) are at the edge of the cluster.

In [Figure 1-1](#), the cluster command switch has ports assigned to VLANs 16 and 62. The CDP hop count is three. The cluster command switch discovers switches 11, 12, 13, and 14 because they are within three hops from the edge of the cluster. It does not discover switch 15 because it is four hops from the edge of the cluster.

Figure 1-1 Discovery Through CDP Hops




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

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

Figure 1-2 shows that the cluster command switch discovers the switch that is connected to a third-party hub. However, the cluster command switch does not discover the switch that is connected to a Catalyst 5000 switch.

Figure 1-2 Discovery Through Non-CDP-Capable and Noncluster-Capable Devices



Discovery Through Different VLANs

If the cluster command switch is a Catalyst 2970, Catalyst 3550, Catalyst 3560, or Catalyst 3750 switch, the cluster can have cluster member switches in different VLANs. As cluster member switches, they must be connected through at least one VLAN in common with the cluster command switch. The cluster command switch in [Figure 1-3](#) has ports assigned to VLANs 9, 16, and 62 and therefore discovers the switches in those VLANs. It does not discover the switch in VLAN 50. It also does not discover the switch in VLAN 16 in the first column because the cluster command switch has no VLAN connectivity to it.


Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL cluster member switches must be connected to the cluster command switch through their management VLAN. For information about discovery through management VLANs, see the [“Discovery Through Different Management VLANs”](#) section on page 1-7. For more information about VLANs, see [Chapter 1, “Configuring VLANs.”](#)



Note

For additional considerations about VLANs in switch stacks, see the [“Switch Clusters and Switch Stacks”](#) section on page 1-13.

Figure 1-3 Discovery Through Different VLANs



Discovery Through Different Management VLANs

Catalyst 2970, Catalyst 3550, Catalyst 3560, or Catalyst 3750 cluster command switches can discover and manage cluster member switches in different VLANs and different management VLANs. As cluster member switches, they must be connected through at least one VLAN in common with the cluster command switch. They do not need to be connected to the cluster command switch through their management VLAN. The default management VLAN is VLAN 1.

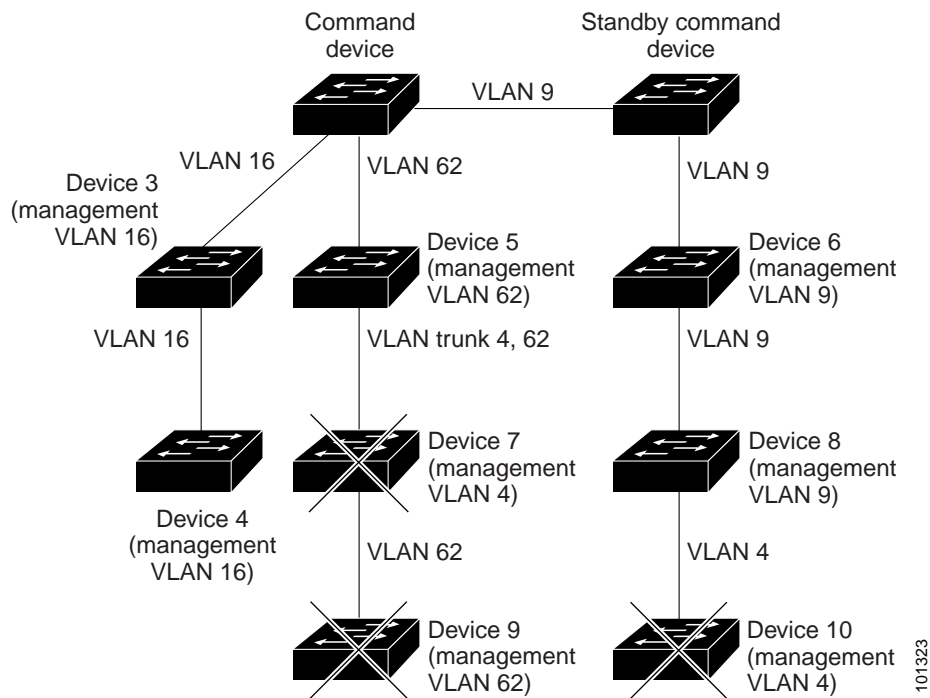
Note

If the switch cluster has a Catalyst 3750 or 2975 switch or has a switch stack, that switch or switch stack must be the cluster command switch.

The cluster command switch and standby command switch in [Figure 1-4](#) (assuming they are Catalyst 2960, Catalyst 2970, Catalyst 2975, Catalyst 3550, Catalyst 3560, or Catalyst 3750 cluster command switches) have ports assigned to VLANs 9, 16, and 62. The management VLAN on the cluster command switch is VLAN 9. Each cluster command switch discovers the switches in the different management VLANs except these:

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

Figure 1-4 Discovery Through Different Management VLANs with a Layer 3 Cluster Command Switch



Discovery of Newly Installed Switches


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

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

The cluster command switch in [Figure 1-5](#) belongs to VLANs 9 and 16. When new cluster-capable switches join the cluster:

- One cluster-capable switch and its access port are assigned to VLAN 9.
- The other cluster-capable switch and its access port are assigned to management VLAN 16.

Figure 1-5 Discovery of Newly Installed Switches



HSRP and Standby Cluster Command Switches

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

- For a cluster command switch stack, a standby cluster command switch is necessary if the entire switch stack fails. However, if only the stack master in the command switch stack fails, the switch stack elects a new stack master and resumes its role as the cluster command switch stack.
- For a cluster command switch that is a standalone switch, configure a standby cluster command switch to take over if the primary cluster command switch fails.

A *cluster standby group* is a group of command-capable switches that meet the requirements described in the “[Standby Cluster Command Switch Characteristics](#)” section on page 1-3. Only one cluster standby group can be assigned per cluster.

The switches in the cluster standby group are ranked according to HSRP priorities. The switch with the highest priority in the group is the *active cluster command switch* (AC). The switch with the next highest priority is the *standby cluster command switch* (SC). The other switches in the cluster standby group are the *passive cluster command switches* (PC). If the active cluster command switch and the standby cluster command switch become disabled *at the same time*, the passive cluster command switch with the highest priority becomes the active cluster command switch. For the limitations to automatic discovery, see the “[Automatic Recovery of Cluster Configuration](#)” section on page 1-11.



Note

The HSRP standby hold time interval should be greater than or equal to three times the hello time interval. The default HSRP standby hold time interval is 10 seconds. The default HSRP standby hello time interval is 3 seconds.

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

- [Virtual IP Addresses, page 1-10](#)
- [Other Considerations for Cluster Standby Groups, page 1-10](#)
- [Automatic Recovery of Cluster Configuration, page 1-11](#)

Virtual IP Addresses

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

If the active cluster command switch fails, the standby cluster command switch assumes ownership of the virtual IP address and becomes the active cluster command switch. The passive switches in the cluster standby group compare their assigned priorities to decide the new standby cluster command switch. The passive standby switch with the highest priority then becomes the standby cluster command switch. When the previously active cluster command switch becomes active again, it resumes its role as the active cluster command switch, and the current active cluster command switch becomes the standby cluster command switch again. For more information about IP address in switch clusters, see the [“IP Addresses” section on page 1-12](#).

Other Considerations for Cluster Standby Groups



Note

For additional considerations about cluster standby groups in switch stacks, see the [“Switch Clusters and Switch Stacks” section on page 1-13](#).

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

These requirements also apply:

- Standby cluster command switches must be the same type of switches as the cluster command switch. For example, if the cluster command switch is a Catalyst 2960 switch, the standby cluster command switches must also be Catalyst 2960 switches. If the cluster command switch is a Catalyst 2960-S switch, the standby cluster command switches must also be Catalyst 2960-S switches. Refer to the switch configuration guide of other cluster-capable switches for their requirements on standby cluster command switches.

If your switch cluster has a Catalyst 2960 switch or a Cisco FlexStack (a stack that contains only 2960-S switches), it should be the cluster command switch.

- Only one cluster standby group can be assigned to a cluster. You can have more than one router-redundancy standby group.

- All standby-group members must be members of the cluster.




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

- Each standby-group member (Figure 1-6) must be connected to the cluster command switch through the same VLAN. In this example, the cluster command switch and standby cluster command switches are Catalyst 2970, Catalyst 3550, Catalyst 3560, or Catalyst 3750 cluster command switches. Each standby-group member must also be redundantly connected to each other through at least one VLAN in common with the switch cluster.

Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL cluster member switches must be connected to the cluster standby group through their management VLANs. For more information about VLANs in switch clusters, see these sections:

- “Discovery Through Different VLANs” section on page 1-7
- “Discovery Through Different Management VLANs” section on page 1-7

Figure 1-6 VLAN Connectivity between Standby-Group Members and Cluster Members



Automatic Recovery of Cluster Configuration

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

Automatic discovery has these limitations:

- This limitation applies only to clusters that have Catalyst 2950, Catalyst 3550, Catalyst 3560, and Catalyst 3750 command and standby cluster command switches: If the active cluster command switch and standby cluster command switch become disabled *at the same time*, the passive cluster command switch with the highest priority becomes the active cluster command switch. However, because it was a passive standby cluster command switch, the previous cluster command switch *did*

not forward cluster-configuration information to it. The active cluster command switch only forwards cluster-configuration information to the standby cluster command switch. You must therefore rebuild the cluster.

- This limitation applies to all clusters: If the active cluster command switch fails and there are more than two switches in the cluster standby group, the new cluster command switch does not discover any Catalyst 1900, Catalyst 2820, and Catalyst 2916M XL cluster member switches. You must re-add these cluster member switches to the cluster.
- This limitation applies to all clusters: If the active cluster command switch fails and becomes active again, it does not discover any Catalyst 1900, Catalyst 2820, and Catalyst 2916M XL cluster member switches. You must again add these cluster member switches to the cluster.

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

IP Addresses

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

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

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

For more information about IP addresses, see [Chapter 1, “Assigning the Switch IP Address and Default Gateway.”](#)

Hostnames

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

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

If a switch has a hostname, it retains that name when it joins a cluster and when it leaves the cluster.

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

Passwords

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

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

For more information about passwords, see the [“Preventing Unauthorized Access to Your Switch” section on page 1-1](#).

For password considerations specific to the Catalyst 1900 and Catalyst 2820 switches, refer to the installation and configuration guides for those switches.

SNMP Community Strings

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

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

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

The switches support an unlimited number of community strings and string lengths. For more information about SNMP and community strings, see [Chapter 1, “Configuring SNMP.”](#)

For SNMP considerations specific to the Catalyst 1900 and Catalyst 2820 switches, refer to the installation and configuration guides specific to those switches.

Switch Clusters and Switch Stacks

A *switch cluster* can have one or more Catalyst 2960-S switch stacks. Each switch stack can act as the cluster command switch or as a single cluster member. [Table 1-2](#) describes the basic differences between switch stacks and switch clusters. For more information about switch stacks, see [Chapter 1, “Managing Switch Stacks.”](#)



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

Table 1-2 Basic Comparison of Switch Stacks and Switch Clusters

Switch Stack	Switch Cluster
Made up of Catalyst 2960-S switches only	Made up of cluster-capable switches, such as Catalyst 3750, Catalyst 3550, and Catalyst 2960-S switches
Stack members are connected through StackWise ports	Cluster members are connected through LAN ports
Requires one <i>stack master</i> and supports up to four other <i>stack members</i>	Requires 1 <i>cluster command switch</i> and supports up to 15 other <i>cluster member switches</i>
Can be a cluster command switch or a cluster member switch	Cannot be a stack master or stack member
Stack master is the single point of <i>complete</i> management for all stack members in a particular switch stack	Cluster command switch is the single point of <i>some</i> management for all cluster members in a particular switch cluster
Back-up stack master is automatically determined in case the stack master fails	Standby cluster command switch must be pre-assigned in case the cluster command switch fails
Switch stack supports up to four simultaneous stack master failures	Switch cluster supports only one cluster command switch failure at a time
Stack members (as a switch stack) behave and is presented as a single, unified system in the network	Cluster members are various, independent switches that are not managed as and do not behave as a unified system
Integrated management of stack members through a single configuration file	Cluster members have separate, individual configuration files
Stack- and interface-level configurations are stored on each stack member	Cluster configuration are stored on the cluster command switch and the standby cluster command switch
New stack members are automatically added to the switch stack	New cluster members must be manually added to the switch cluster

Recall that stack members work together to behave as a unified system (as a single switch stack) in the network and are presented to the network as such by Layer 2 and Layer 3 protocols. Therefore, the switch cluster recognizes switch stacks, not individual stack members, as eligible cluster members. Individual stack members cannot join a switch cluster or participate as separate cluster members. Because a switch cluster must have 1 cluster command switch and can have up to 15 cluster members, a cluster can potentially have up to 16 switch stacks, totalling 144 devices.

Cluster configuration of switch stacks is through the stack master.

These are considerations to keep in mind when you have switch stacks in switch clusters:

- If the cluster command switch is not a Catalyst 2960-S switch or switch stack and a new stack master is elected in a cluster member switch stack, the switch stack loses its connectivity to the switch cluster if there are no redundant connections between the switch stack and the cluster command switch. You must add the switch stack to the switch cluster.
- If the cluster command switch is a switch stack and new stack masters are simultaneously elected in the cluster command switch stack and in cluster member switch stacks, connectivity between the switch stacks is lost if there are no redundant connections between the switch stack and the cluster command switch. You must add the switch stacks to the cluster, including the cluster command switch stack.

- All stack members should have redundant connectivity to all VLANs in the switch cluster. Otherwise, if a new stack master is elected, stack members connected to any VLANs not configured on the new stack master lose their connectivity to the switch cluster. You must change the VLAN configuration of the stack master or the stack members and add the stack members back to the switch cluster.
- If a cluster member switch stack reloads and a new stack master is elected, the switch stack loses connectivity with the cluster command switch. You must add the switch stack back to the switch cluster.
- If a cluster command switch stack reloads, and the original stack master is not re-elected, you must rebuild the entire switch cluster.

For more information about switch stacks, see [Chapter 1, “Managing Switch Stacks,”](#)

TACACS+ and RADIUS

If Terminal Access Controller Access Control System Plus (TACACS+) is configured on a cluster member, it must be configured on all cluster members. Similarly, if RADIUS is configured on a cluster member, it must be configured on all cluster members. Further, the same switch cluster cannot have some members configured with TACACS+ and other members configured with RADIUS.

For more information about TACACS+, see the [“Controlling Switch Access with TACACS+”](#) section on page 1-10. For more information about RADIUS, see the [“Controlling Switch Access with RADIUS”](#) section on page 1-17.

LRE Profiles

A configuration conflict occurs if a switch cluster has Long-Reach Ethernet (LRE) switches that use both private and public profiles. If one LRE switch in a cluster is assigned a public profile, all LRE switches in that cluster must have that same public profile. Before you add an LRE switch to a cluster, make sure that you assign it the same public profile used by other LRE switches in the cluster.

A cluster can have a mix of LRE switches that use different private profiles.

Using the CLI to Manage Switch Clusters

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

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

```
switch# rcommand 3
```

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

The Telnet session accesses the member-switch CLI at the same privilege level as on the cluster command switch. The Cisco IOS commands then operate as usual. For instructions on configuring the switch for a Telnet session, see the [“Disabling Password Recovery”](#) section on page 1-5.

**Note**

The CLI supports creating and maintaining switch clusters with up to 16 switch stacks. For more information about switch stack and switch cluster, see the [“Switch Clusters and Switch Stacks” section on page 1-13](#).

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

Catalyst 1900 and Catalyst 2820 CLI Considerations

If your switch cluster has Catalyst 1900 and Catalyst 2820 switches running standard edition software, the Telnet session accesses the management console (a menu-driven interface) if the cluster command switch is at privilege level 15. If the cluster command switch is at privilege level 1 to 14, you are prompted for the password to access the menu console.

Command-switch privilege levels map to the Catalyst 1900 and Catalyst 2820 cluster member switches running standard and Enterprise Edition Software as follows:

- If the command-switch privilege level is 1 to 14, the cluster member switch is accessed at privilege level 1.
- If the command-switch privilege level is 15, the cluster member switch is accessed at privilege level 15.

**Note**

The Catalyst 1900 and Catalyst 2820 CLI is available only on switches running Enterprise Edition Software.

For more information about the Catalyst 1900 and Catalyst 2820 switches, refer to the installation and configuration guides for those switches.

Using SNMP to Manage Switch Clusters

When you first power on the switch, SNMP is enabled if you enter the IP information by using the setup program and accept its proposed configuration. If you did not use the setup program to enter the IP information and SNMP was not enabled, you can enable it as described in the [“Configuring SNMP” section on page 1-6](#). On Catalyst 1900 and Catalyst 2820 switches, SNMP is enabled by default.

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


**Note**

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

If the cluster member switch does not have an IP address, the cluster command switch redirects traps from the cluster member switch to the management station, as shown in [Figure 1-7](#). If a cluster member switch has its own IP address and community strings, the cluster member switch can send traps directly to the management station, without going through the cluster command switch.

If a cluster member switch has its own IP address and community strings, they can be used in addition to the access provided by the cluster command switch. For more information about SNMP and community strings, see [Chapter 1, “Configuring SNMP.”](#)

Figure 1-7 *SNMP Management for a Cluster*





CHAPTER 1

Managing Switch Stacks

This chapter provides the concepts and procedures to manage Catalyst 2960-S stacks, also referred to as Cisco FlexStacks. See the command reference for command syntax and usage information.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

- [Understanding Stacks, page 1-1](#)
- [Configuring the Switch Stack, page 1-16](#)
- [Accessing the CLI of a Specific Member, page 1-21](#)
- [Displaying Stack Information, page 1-22](#)
- [Troubleshooting Stacks, page 1-22](#)

For other switch stack-related information, such as cabling the switches through their stack ports and using the LEDs for switch stack status, see the hardware installation guide.

Understanding Stacks

A *switch stack* is a set of up to four Catalyst 2960-S switches connected through their stack ports. One of the switches controls the operation of the stack and is called the *stack master*. The stack master and the other switches in the stack are *stack members*. Layer 2 protocol presents the entire switch stack as a single entity to the network.



Note

A switch stack is different from a *switch cluster*. A switch cluster is a set of switches connected through their LAN ports, such as the 10/100/1000 ports. For more information about how switch stacks differ from switch clusters, see the “Planning and Creating Clusters” chapter in the *Getting Started with Cisco Network Assistant* on Cisco.com.

The master is the single point of stack-wide management. From the master, you configure:

- System-level (global) features that apply to all members
- Interface-level features for each member

If the stack master is running the cryptographic version (that is, supports encryption) of the software, the encryption features are available.

Every member is uniquely identified by its own *stack member number*.

All members are eligible masters. If the master becomes unavailable, the remaining members elect a new master from among themselves. One of the factors is the *stack member priority value*. The switch with the highest stack-member priority-value becomes the master.

The system-level features supported on the master are supported on the entire stack.

The master contains the saved and running configuration files for the stack. The configuration files include the system-level settings for the stack and the interface-level settings for each member. Each member has a current copy of these files for back-up purposes.

You manage the stack through a single IP address. The IP address is a system-level setting and is not specific to the master or to any other member. You can manage the stack through the same IP address even if you remove the master or any other member from the stack.

You can use these methods to manage stacks:

- Network Assistant (available on Cisco.com)
- Command-line interface (CLI) over a serial connection to the console port of any member
- A network management application through the Simple Network Management Protocol (SNMP)



Note Use SNMP to manage network features across the stack that are defined by supported MIBs. The switch does not support MIBs to manage stacking-specific features such as stack membership and election.

- CiscoWorks network management software

To manage stacks, you should understand:

- These concepts on stack formations:
 - [Stack Membership, page 1-3](#)
 - [Master Election, page 1-5](#)
- These concepts on stack and member configurations:
 - [Stack MAC Address, page 1-6](#)
 - [Member Numbers, page 1-6](#)
 - [Member Priority Values, page 1-7](#)
 - [Stack Offline Configuration, page 1-7](#)
 - [Stack Software Compatibility Recommendations, page 1-9](#)
 - [Stack Protocol Version Compatibility, page 1-9](#)
 - [Major Version Number Incompatibility Among Switches, page 1-9](#)
 - [Minor Version Number Incompatibility Among Switches, page 1-10](#)
 - [Incompatible Software and Member Image Upgrades, page 1-13](#)
 - [Stack Configuration Files, page 1-13](#)
 - [Additional Considerations for System-Wide Configuration on Switch Stacks, page 1-13](#)
 - [Stack Management Connectivity, page 1-14](#)
 - [Stack Configuration Scenarios, page 1-15](#)
- This concept on stack topology changes:
 - [Data Recovery After Stack Topology Changes, page 1-16](#)

Stack Membership

**Note**

A switch stack can have only Catalyst 2960-S stack members.

A *standalone switch* is a stack with one member that is also the master. You can connect one standalone switch to another (Figure 1-1 on page 1-4) to create a stack containing two stack members, with one of them as the master. You can connect standalone switches to an existing stack (Figure 1-2 on page 1-4) to increase the stack membership.

If you replace a stack member with an identical model, the new switch functions with the same configuration as the replaced switch (assuming that the new switch is using the same member number as the replaced switch). For information about the benefits of provisioning a switch stack, see the “[Stack Offline Configuration](#)” section on page 1-7. For information about replacing a failed switch, see the “[Troubleshooting](#)” chapter in the hardware installation guide.

The operation of the stack continues uninterrupted during membership changes unless you remove the master or you add powered-on standalone switches or stacks.


**Note**

To prevent interrupted stack operations, make sure the switches that you add to or remove from the stack are powered off.

After adding or removing members, make sure that the stack ring is operating at full bandwidth (20 Gb/s). Press the Mode button on a member until the Stack mode LED is on. The last two port LEDs on all switches in the stack should be green. If any one or both of any the last two port LEDs are not green, the stack is not operating at full bandwidth.

- Adding powered-on switches (merging) causes the masters of the merging stacks to elect a master from among themselves. The new master keeps its role and configuration and so do its members. All remaining switches, including the former masters, reload and join the stack as members. They change their member numbers to the lowest available numbers and use the configuration of the new master.
- Removing powered-on members divides (partitions) the stack into two or more switch stacks, each with the same configuration. This can create an IP address configuration conflict in your network. If you want the stacks to remain separate, change the IP address or addresses of the newly created stacks.

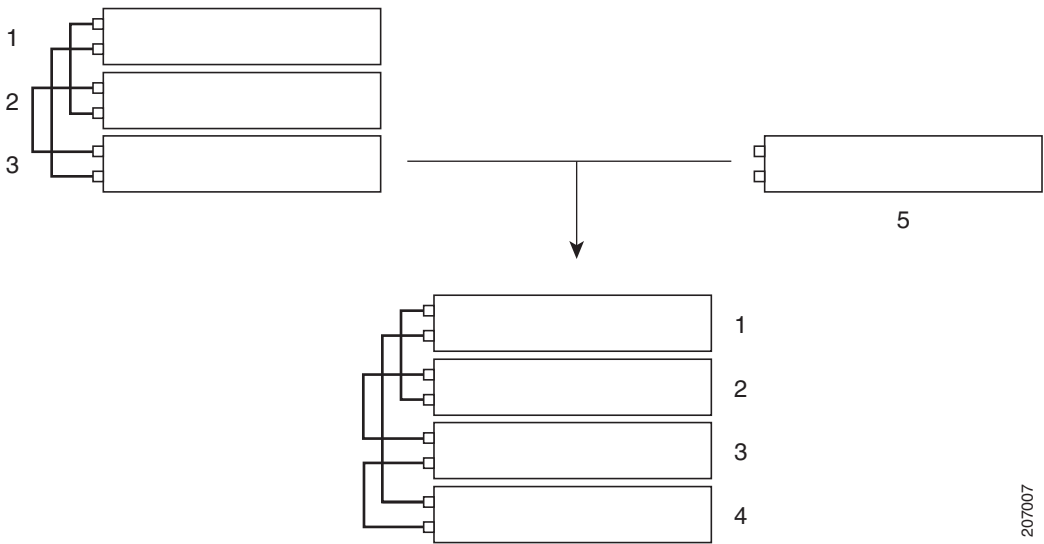
Figure 1-1 *Creating a Switch Stack from Two Standalone Switches*



207006

1	Standalone switch	3	Stack member 2 and stack master
2	Stack member 1		

Figure 1-2 *Adding a Standalone Switch to a Switch Stack*



207007

1	Stack member 1	4	Stack member 4
2	Stack member 2 and stack master	5	Standalone switch
3	Stack member 3		

For information about cabling and powering switch stacks, see the “Switch Installation” chapter in the hardware installation guide.

Master Election

The stack master is elected based on one of these factors in the order listed:

1. The switch that is currently the stack master.
2. The switch with the highest stack member priority value.



Note We recommend you assign the highest priority value to the switch that you want to be the master. The switch is then re-elected as master if a re-election occurs.

3. The switch that has the configuration file.
4. The switch with the highest uptime.
5. The switch with the lowest MAC address.

A stack master keeps its role unless one of these events occurs:

- The stack is reset.*
- The master is removed from the stack.
- The master is reset or powered off.
- The master fails.
- The stack membership is increased by adding powered-on standalone switches or switch stacks.*

In the events marked by an asterisk (*), the current stack master *might* be re-elected based on the listed factors.

When you power on or reset an entire stack, some stack members *might not* participate in the master election.

- All members participate in re-elections.
- Members that are powered on within the same 20-second time frame participate in the master election and have a chance to become the master.
- Members that are powered on after the 20-second time frame do not participate in this initial election and only become members.

The new master is available after a few seconds. In the meantime, the switch stack uses the forwarding tables in memory to minimize network disruption. The physical interfaces on the other available stack members are not affected while a new stack master is elected and is resetting.

When a new master is elected and the previous stack master becomes available, the previous master *does not* resume its role as stack master.

For all powering considerations that affect stack-master elections, see the “Switch Installation” chapter in the hardware installation guide.

Stack MAC Address

The MAC address of the master determines the stack MAC address.

When the stack initializes, the MAC address of the master determines the bridge ID that identifies the stack in the network.

If the master changes, the MAC address of the *new* master determines the new bridge ID. However, when the persistent MAC address feature is enabled, there is an approximate 4-minute delay before the stack MAC address changes. During this time period, if the previous master rejoins the stack, the stack continues to use its MAC address as the stack MAC address, even if the switch is now a member and not a master. If the previous master does not rejoin the stack during this period, the stack takes the MAC address of the new stack master as the stack MAC address. See the [“Enabling Persistent MAC Address” section on page 1-17](#) for more information.

Member Numbers

The member number (1 to 4) identifies each member in the stack. The member number also determines the interface-level configuration that a member uses.

A new, out-of-the-box switch (one that has not joined a stack or has not been manually assigned a member number) ships with a default member number of 1. When it joins a stack, its default stack member number changes to the lowest available member number in the stack.

Members in the same stack cannot have the same member number.

- If you manually change the member number by using the **switch** *current-stack-member-number renumber new-stack-member-number* global configuration command, the new number goes into effect after that member resets (or after you use the **reload slot** *stack-member-number* privileged EXEC command) and only if that number is not already changed.

You can also change the stack member number is by using the SWITCH_NUMBER environment variable.

If the number is being used by another member in the stack, the switch selects the lowest available number in the stack.

If you manually change the member number and no interface-level configuration is associated with that number, that member resets to its default configuration.

You cannot use the **switch** *current-stack-member-number renumber new-stack-member-number* global configuration command on a provisioned switch. If you do, the command is rejected.

- If you move a stack member to a different switch stack, the stack member keeps its number only if the number is not being used by another member in the stack. If it is being used by another member in the stack, the switch selects the lowest available number in the stack.

See the following sections for information about stack member configuration:

- The procedure to change a member number, see the [“Assigning a Member Number” section on page 1-19](#).
- The SWITCH_NUMBER environment variable, see the [“Controlling Environment Variables” section on page 1-20](#).
- Member numbers and configurations, see the [“Stack Configuration Files” section on page 1-13](#).
- Merging stacks, see the [“Stack Membership” section on page 1-3](#).

Member Priority Values

A high priority value for a member increases the chance that it will be elected master and keep its member number. The priority value can be 1 to 15. The default priority value is 1.



Note

We recommend that you assign the highest priority value to the switch that you want to be the stack master. The switch is then re-elected as master if a re-election occurs.

The new priority value takes effect immediately but does not affect the current master until the current master or the stack resets.

Stack Offline Configuration

You can use the offline configuration feature to *provision* (to configure) a new switch before it joins the stack. You can configure the member number, the switch type, and the interfaces associated with a switch that is not yet part of the stack. That configuration is the *provisioned configuration*. The switch to be added to the stack and to get this configuration is the *provisioned switch*.

The provisioned configuration is automatically created when a switch is added to a stack and when no provisioned configuration exists. You can manually create the provisioned configuration by using the **switch stack-member-number provision type** global configuration command.

When you configure the interfaces for a provisioned switch (for example, as part of a VLAN), the information appears in the stack running configuration whether or not the provisioned switch is part of the stack. The interface for the provisioned switch is not active and does not appear in the display of a specific feature (for example, in the **show vlan** user EXEC command output). Entering the **no shutdown** interface configuration command has no effect.

The startup configuration file ensures that the stack can reload and can use the saved information whether or not the provisioned switch is part of the stack.

Effects of Adding a Provisioned Switch to a Stack

When you add a provisioned switch to the switch stack, the stack applies either the provisioned configuration or the default configuration to it. [Table 1-1](#) lists the events that occur when the switch stack compares the provisioned configuration with the provisioned switch.

Table 1-1 Results of Comparing the Provisioned Configuration with the Provisioned Switch

Scenario		Result
The stack member numbers and the switch types match.	<ol style="list-style-type: none"> 1. If the stack member number of the provisioned switch matches the stack member number in the provisioned configuration on the stack, and 2. If the switch type of the provisioned switch matches the switch type in the provisioned configuration on the stack. 	The switch stack applies the provisioned configuration to the provisioned switch and adds it to the stack.

Table 1-1 Results of Comparing the Provisioned Configuration with the Provisioned Switch (continued)

Scenario		Result
The stack member numbers match but the switch types do not match.	<ol style="list-style-type: none"> 1. If the stack member number of the provisioned switch matches the stack member number in the provisioned configuration on the stack, but 2. The switch type of the provisioned switch does not match the switch type in the provisioned configuration on the stack. 	<p>The switch stack applies the default configuration to the provisioned switch and adds it to the stack.</p> <p>The provisioned configuration is changed to reflect the new information.</p>
The stack member number is not found in the provisioned configuration.		<p>The switch stack applies the default configuration to the provisioned switch and adds it to the stack.</p> <p>The provisioned configuration is changed to reflect the new information.</p>
The stack member number of the provisioned switch is in conflict with an existing stack member.	<p>The stack master assigns a new stack member number to the provisioned switch.</p> <p>The stack member numbers and the switch types match:</p> <ol style="list-style-type: none"> 1. If the new stack member number of the provisioned switch matches the stack member number in the provisioned configuration on the stack, and 2. If the switch type of the provisioned switch matches the switch type in the provisioned configuration on the stack. 	<p>The switch stack applies the provisioned configuration to the provisioned switch and adds it to the stack.</p> <p>The provisioned configuration is changed to reflect the new information.</p>
	<p>The stack member numbers match, but the switch types do not match:</p> <ol style="list-style-type: none"> 1. If the stack member number of the provisioned switch matches the stack member number in the provisioned configuration on the stack, but 2. The switch type of the provisioned switch does not match the switch type in the provisioned configuration on the stack. 	<p>The switch stack applies the default configuration to the provisioned switch and adds it to the stack.</p> <p>The provisioned configuration is changed to reflect the new information.</p>
The stack member number of the provisioned switch is not found in the provisioned configuration.		The switch stack applies the default configuration to the provisioned switch and adds it to the stack.

If you add a provisioned switch that is a different type than specified in the provisioned configuration to a powered-down switch stack and then apply power, the switch stack rejects the (now incorrect) **switch stack-member-number provision type** global configuration command in the startup configuration file. However, during stack initialization, the nondefault interface configuration information in the startup configuration file for the provisioned interfaces (potentially of the wrong type) are executed. Depending on how different the actual switch type is from the previously provisioned switch type, some commands are rejected, and some commands are accepted.

**Note**

If the switch stack does not contain a provisioned configuration for a new switch, the switch joins the stack with the default interface configuration. The switch stack then adds to its running configuration a **switch** *stack-member-number* **provision** *type* global configuration command that matches the new switch.

For configuration information, see the [“Provisioning a New Member for a Stack”](#) section on page 1-20.

Effects of Replacing a Provisioned Switch in a Stack

When a provisioned switch in a switch stack fails, is removed from the stack, and is replaced with another switch, the stack applies either the provisioned configuration or the default configuration to it. The events that occur when the switch stack compares the provisioned configuration with the provisioned switch are the same as those described in the [“Effects of Adding a Provisioned Switch to a Stack”](#) section on page 1-7.

Effects of Removing a Provisioned Switch from a Stack

If you remove a provisioned switch from the switch stack, the configuration associated with the removed stack member remains in the running configuration as provisioned information. To completely remove the configuration, use the **no switch** *stack-member-number* **provision** global configuration command.

Stack Software Compatibility Recommendations

All stack members must run the same Cisco IOS software version to ensure compatibility in the stack protocol version among the members.

Stack Protocol Version Compatibility

The stack protocol version has a *major* version number and a *minor* version number (for example 1.4, where 1 is the major version number and 4 is the minor version number).

Switches with the same Cisco IOS software version have the same stack protocol version. All features function properly across the stack. These switches with the same software version as the master immediately join the stack.

If an incompatibility exists, a system message describes the cause of the incompatibility on the specific stack members. The master sends the message to all members.

For more information, see the [“Major Version Number Incompatibility Among Switches”](#) procedure on page 1-9 and the [“Minor Version Number Incompatibility Among Switches”](#) procedure on page 1-10.

Major Version Number Incompatibility Among Switches

Switches with different Cisco IOS software versions likely have different stack protocol versions. Switches with different major version numbers are incompatible and cannot exist in the same stack.

Minor Version Number Incompatibility Among Switches

Switches with the same major version number but with a different minor version number as the master are considered partially compatible. When connected to a stack, a partially compatible switch enters version-mismatch mode and cannot join the stack as a fully functioning member. The software detects the mismatched software and tries to upgrade (or downgrade) the switch in version-mismatch mode with the stack image or with a tar file image from the stack flash memory. The software uses the automatic upgrade (auto-upgrade) and the automatic advise (auto-advise) features.

The port LEDs on switches in version-mismatch mode will also stay off. Pressing the Mode button does not change the LED mode.

**Note**

Auto-advise and auto-copy identify which images are running by examining the info file and by searching the directory structure on the switch stack. If you download your image by using the **copy tftp:** command instead of by using the **archive download-sw** privileged EXEC command, the correct directory structure is not properly created. For more information about the info file, see the “[tar File Format of Images on a Server or Cisco.com](#)” section on page 1-25.

Understanding Auto-Upgrade and Auto-Advise

When the software detects mismatched software and tries to upgrade the switch in version-mismatch mode, two software processes are involved: automatic upgrade and automatic advise.

- The automatic upgrade (auto-upgrade) process includes an auto-copy process and an auto-extract process. By default, auto-upgrade is enabled (the **boot auto-copy-sw** global configuration command is enabled). You can disable auto-upgrade by using the **no boot auto-copy-sw** global configuration command on the master. You can check the status of auto-upgrade by using the **show boot** privileged EXEC command and by checking the *Auto upgrade* line in the display.
 - Auto-copy automatically copies the software image running on any member to the switch in version-mismatch mode to upgrade (auto-upgrade) it. Auto-copy occurs if auto-upgrade is enabled, if there is enough flash memory in the switch in version-mismatch mode, and if the software image running on the stack is suitable for the switch in version-mismatch mode.

**Note**

A switch in version-mismatch mode might not run all released software. For example, new switch hardware is not recognized in earlier versions of software.

- Automatic extraction (auto-extract) occurs when the auto-upgrade process cannot find the appropriate software in the stack to copy to the switch in version-mismatch mode. In that case, the auto-extract process searches all switches in the stack, whether they are in version-mismatch mode or not, for the tar file needed to upgrade the switch stack or the switch in version-mismatch mode. The tar file can be in any flash file system in the stack (including the switch in version-mismatch mode). If a tar file suitable for the switch in version-mismatch mode is found, the process extracts the file and automatically upgrades that switch.

The auto-upgrade (auto-copy and auto-extract) processes start a few minutes after the mismatched software is detected.

When the auto-upgrade process is complete, the switch that was in version-mismatch mode reloads and joins the stack as a fully functioning member. If you have both stack cables connected during the reload, network downtime does not occur because the stack operates on two rings.

- Automatic advise (auto-advise)—when the auto-upgrade process cannot find appropriate version-mismatch member software to copy to the switch in version-mismatch mode, the auto-advise process tells you the command (**archive copy-sw** or **archive download-sw** privileged EXEC command) and the image name (tar filename) needed to manually upgrade the switch stack or the switch in version-mismatch mode. The recommended image can be the running stack image or a tar file in any flash file system in the stack (including the switch in version-mismatch mode). If an appropriate image is not found in the stack flash file systems, the auto-advise process tells you to install new software on the stack. Auto-advise cannot be disabled, and there is no command to check its status.

The auto-advise software does *not* give suggestions when the stack software and the software of the switch in version-mismatch mode do not contain the same feature sets. The same events occur when cryptographic and noncryptographic images are running.

You can use the **archive-download-sw /allow-feature-upgrade** privileged EXEC command to allow installing an image with a different feature set.

Auto-Upgrade and Auto-Advise Example Messages

When you add a switch that has a different minor version number to the stack, the software displays messages in sequence (assuming that there are no other system messages generated by the switch).

This example shows that the stack detected a new switch that is running a different minor version number than the stack. Auto-copy launches, finds suitable software to copy from a member to the switch in version-mismatch mode, upgrades the switch in version-mismatch mode, and then reloads it:

```
*Mar 11 20:31:19.247:%STACKMGR-6-STACK_LINK_CHANGE:Stack Port 2 Switch 2 has changed to
state UP
*Mar 11 20:31:23.232:%STACKMGR-6-SWITCH_ADDED_VM:Switch 1 has been ADDED to the stack
(VERSION_MISMATCH)
*Mar 11 20:31:23.291:%STACKMGR-6-SWITCH_ADDED_VM:Switch 1 has been ADDED to the stack
(VERSION_MISMATCH) (Stack_1-3)
*Mar 11 20:33:23.248:%IMAGEMGR-6-AUTO_COPY_SW_INITIATED:Auto-copy-software process
initiated for switch number(s) 1
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:Searching for stack member to act
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:as software donor...
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:Found donor (system #2) for
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:member(s) 1
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:System software to be uploaded:
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:System Type:          0x00000000
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:archiving (directory)
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:archiving /.bin (4945851 bytes)
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:archiving /info (450 bytes)
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:archiving info (104 bytes)
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:examining image...
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:extracting info (104 bytes)
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:extracting /info (450 bytes)
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:extracting info (104 bytes)
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:Stacking Version Number:1.4
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:System Type:          0x00000000
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW: Ios Image File Size:  0x004BA200
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW: Total Image File Size:0x00818A00
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW: Minimum Dram required:0x08000000
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW: Image Suffix:universalk9-122-53.SE
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW: Image Directory:
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW: Image Name:
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW: Image
Feature: IP|LAYER_3|PLUS|MIN_DRAM_MEG=128
```

```

*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:Old image for switch 1:flash1:
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW: Old image will be deleted after download.
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:Extracting images from archive into flash on
switch 1...
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW: (directory)
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:extracting / (4945851 bytes)
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:extracting /info (450 bytes)
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:extracting info (104 bytes)
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:Installing (renaming):`flash1:update/' ->
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:      `flash1:'
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:New software image installed in flash1:
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:Removing old image:flash1:
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:All software images installed.
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:Requested system reload in progress...
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:Software successfully copied to
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:system(s) 1
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:Done copying software
*Mar 11 20:36:15.038:%IMAGEMGR-6-AUTO_COPY_SW:Reloading system(s) 1

```

This example shows that the stack detected a new switch that is running a different minor version number than the stack. Auto-copy launches but cannot find software in the stack to copy to the switch in version-mismatch mode to make it compatible with the stack. The auto-advise process launches and recommends that you download a tar file from the network to the switch in version-mismatch mode:

```

*Mar 1 00:01:11.319:%STACKMGR-6-STACK_LINK_CHANGE:Stack Port 2 Switch 2 has changed to
state UP
*Mar 1 00:01:15.547:%STACKMGR-6-SWITCH_ADDED_VM:Switch 1 has been ADDED to the stack
(VERSION_MISMATCH)
stack_2#
*Mar 1 00:03:15.554:%IMAGEMGR-6-AUTO_COPY_SW_INITIATED:Auto-copy-software process
initiated for switch number(s) 1
*Mar 1 00:03:15.554:%IMAGEMGR-6-AUTO_COPY_SW:
*Mar 1 00:03:15.554:%IMAGEMGR-6-AUTO_COPY_SW:Searching for stack member to act
*Mar 1 00:03:15.554:%IMAGEMGR-6-AUTO_COPY_SW:as software donor...
*Mar 1 00:03:15.554:%IMAGEMGR-6-AUTO_COPY_SW:Software was not copied
*Mar 1 00:03:15.562:%IMAGEMGR-6-AUTO_ADVISE_SW_INITIATED:Auto-advise-software process
initiated for switch number(s) 1
*Mar 1 00:04:22.537:%IMAGEMGR-6-AUTO_ADVISE_SW:
*Mar 1 00:04:22.537:%IMAGEMGR-6-AUTO_ADVISE_SW:
*Mar 1 00:04:22.537:%IMAGEMGR-6-AUTO_ADVISE_SW:Systems with incompatible software
*Mar 1 00:04:22.537:%IMAGEMGR-6-AUTO_ADVISE_SW:have been added to the stack. The
*Mar 1 00:04:22.537:%IMAGEMGR-6-AUTO_ADVISE_SW:storage devices on all of the stack
*Mar 1 00:04:22.537:%IMAGEMGR-6-AUTO_ADVISE_SW:members have been scanned, and it has
*Mar 1 00:04:22.537:%IMAGEMGR-6-AUTO_ADVISE_SW:been determined that the stack can be
*Mar 1 00:04:22.537:%IMAGEMGR-6-AUTO_ADVISE_SW:repaired by issuing the following
*Mar 1 00:04:22.537:%IMAGEMGR-6-AUTO_ADVISE_SW:command(s) :
*Mar 1 00:04:22.537:%IMAGEMGR-6-AUTO_ADVISE_SW:
*Mar 1 00:04:22.537:%IMAGEMGR-6-AUTO_ADVISE_SW:      archive download-sw /force-reload
/overwrite /dest 1 flash1:.tar
*Mar 1 00:04:22.537:%IMAGEMGR-6-AUTO_ADVISE_SW:

```

For information about using the **archive download-sw** privileged EXEC command, see the [“Working with Software Images”](#) section on page 1-24.

Incompatible Software and Member Image Upgrades

You can upgrade a switch that has an incompatible software image by using the **archive copy-sw** privileged EXEC command to copy the software image from an existing member. That switch automatically reloads with the new image and joins the stack as a fully functioning member.

For more information, see the [“Copying an Image File from One Stack Member to Another”](#) section on page 1-38.

Stack Configuration Files

The master has the saved and running configuration files for the stack. All members periodically receive synchronized copies of the configuration files from the master. If the master becomes unavailable, any member assuming the role of master has the latest configuration files.

- System-level (global) configuration settings—such as IP, STP, VLAN, and SNMP settings—that apply to all members
- Member interface-specific configuration settings, which are specific for each member

A new, out-of-box switch joining a stack uses the system-level settings of that stack. If a switch is moved to a different stack, it loses its saved configuration file and uses the system-level configuration of the new stack.

The interface-specific configuration of each member is associated with its member number. A stack member keeps its number unless it is manually changed or it is already used by another member in the same stack.

- If an interface-specific configuration does not exist for that member number, the member uses its default interface-specific configuration.
- If an interface-specific configuration exists for that member number, the member uses the interface-specific configuration associated with that member number.

If you replace a failed member with an identical model, the replacement member automatically uses the same interface-specific configuration. You do not need to reconfigure the interface settings. The replacement switch must have the same member number as the failed switch.

You back up and restore the stack configuration in the same way as you do for a standalone switch configuration.

For information about

- The benefits of provisioning a switch stack, see the [“Stack Offline Configuration”](#) section on page 1-7.
- File systems and configuration files, see [Appendix 1, “Working with the Cisco IOS File System, Configuration Files, and Software Images.”](#)

Additional Considerations for System-Wide Configuration on Switch Stacks

- “Planning and Creating Clusters” chapter in the *Getting Started with Cisco Network Assistant*, available on Cisco.com
- [“MAC Addresses and Switch Stacks”](#) section on page 1-15
- [“802.1x Authentication and Switch Stacks”](#) section on page 1-11

- “VTP and Switch Stacks” section on page 1-9
- “Spanning Tree and Switch Stacks” section on page 1-12
- “MSTP and Switch Stacks” section on page 1-8
- “DHCP Snooping and Switch Stacks” section on page 1-7
- “IGMP Snooping and Switch Stacks” section on page 1-6
- “Port Security and Switch Stacks” section on page 1-18
- “CDP and Switch Stacks” section on page 1-2
- “SPAN and RSPAN and Switch Stacks” section on page 1-10
- “Configuring QoS” section on page 1-1
- “ACLs and Switch Stacks” section on page 1-5
- “EtherChannel and Switch Stacks” section on page 1-10
- “IPv6 and Switch Stacks” section on page 1-6

Stack Management Connectivity

You manage the stack and the member interfaces through the master. You can use the CLI, SNMP, Network Assistant, and CiscoWorks network management applications. You cannot manage members as individual switches.

- [Stack Through an IP Address, page 1-14](#)
- [Stack Through an SSH Session, page 1-14](#)
- [Stack Through Console Ports, page 1-15](#)
- [Specific Members, page 1-15](#)

Stack Through an IP Address

The stack is managed through a system-level IP address. You can still manage the stack through the same IP address even if you remove the master or any other stack member from the stack, provided there is IP connectivity.

**Note**

Members keep their IP addresses when you remove them from a stack. To avoid having two devices with the same IP address in your network, change the IP address of the switch that you removed from the stack.

For related information about switch stack configurations, see the [“Stack Configuration Files” section on page 1-13](#).

Stack Through an SSH Session

The Secure Shell (SSH) connectivity to the stack can be lost if a master running the cryptographic version fails and is replaced by a switch that is running a noncryptographic version. We recommend that a switch running the cryptographic version of the software be the master.

Stack Through Console Ports

You can connect to the master through the console port of one or more members.

Be careful when using multiple CLI sessions to the master. Commands that you enter in one session are not displayed in the other sessions. Therefore, it is possible that you might not be able to identify the session from which you entered a command.

We recommend that you use only one CLI session when managing the stack.

Specific Members

If you want to configure a specific member port, you must include the stack member number in the CLI notation.

To access a specific member, see the [“Accessing the CLI of a Specific Member”](#) section on page 1-21.

Stack Configuration Scenarios

Most of the scenarios in [Table 1-2](#) assume at least two switches are connected through their stack ports.

Table 1-2 *Switch Stack Configuration Scenarios*

Scenario		Result
Master election specifically determined by existing masters	Connect two powered-on stacks through the stack ports.	Only one of the two masters becomes the new stack master.
Master election specifically determined by the member priority value	<ol style="list-style-type: none"> 1. Connect two switches through their stack ports. 2. Use the switch stack-member-number priority new-priority-number global configuration command to set one member with a higher member priority value. 3. Restart both members at the same time. 	The member with the higher priority value is elected master.
Master election specifically determined by the configuration file	<p>Assuming that both members have the same priority value:</p> <ol style="list-style-type: none"> 1. Make sure that one member has a default configuration and that the other member has a saved (nondefault) configuration file. 2. Restart both members at the same time. 	The member with the saved configuration file is elected master.
Master election specifically determined by the MAC address	Assuming that both members have the same priority value, configuration file, and software image, restart both stack members at the same time.	The member with the lower MAC address is elected master.

Table 1-2 Switch Stack Configuration Scenarios (continued)

Scenario		Result
Member number conflict	<p>Assuming that one member has a higher priority value than the other member:</p> <ol style="list-style-type: none"> 1. Ensure that both members have the same member number. If necessary, use the switch current-stack-member-number renumber new-stack-member-number global configuration command. 2. Restart both members at the same time. 	The member with the higher priority value keeps its member number. The other member has a new stack member number.
Add a member	<ol style="list-style-type: none"> 1. Power off the new switch. 2. Through their stack ports, connect the new switch to a powered-on stack. 3. Power on the new switch. 	The master is kept. The new switch is added to the stack.
Master failure	Remove (or power off) the master.	One of the remaining stack members becomes the new master. All other members in the stack remain members and do not restart.
Add more than four members	<ol style="list-style-type: none"> 1. Through their stack ports, connect ten switches. 2. Power on all switches. 	<p>Two switches become masters. One master has four stack members. The other master remains a standalone switch.</p> <p>Use the Mode button and port LEDs on the switches to identify which switches are masters and which switches belong to each master. For information about the Mode button and the LEDs, see the hardware installation guide.</p>

Data Recovery After Stack Topology Changes

When you add or remove a stack member, the stack topology changes. Cisco IOS recovers the data flow.

Configuring the Switch Stack

- [Default Switch Stack Configuration, page 1-17](#)
- [Enabling Persistent MAC Address, page 1-17](#)
- [Assigning Stack Member Information, page 1-19](#)
- [Changing the Stack Membership, page 1-21](#)

Default Switch Stack Configuration

Table 1-3 shows the default switch stack configuration.

Table 1-3 Default Switch Stack Configuration

Feature	Default Setting
Stack MAC address timer	Disabled.
Member number	1
Member priority value	1
Offline configuration	The switch stack is not provisioned.
Persistent MAC address	Disabled.

Enabling Persistent MAC Address

The MAC address of the master determines the stack MAC address. When a master is removed from the stack and a new master takes over, the MAC address of the new master becomes the new stack MAC address. However, you can set the persistent MAC address feature with a time delay before the stack MAC address changes. During this time period, if the previous master rejoins the stack, the stack continues to use that MAC address as the stack MAC address, even if the switch is now a member and not a master. You can also configure stack MAC persistency so that the stack MAC address never changes to the new master MAC address.



Caution

When you configure this feature, a warning message displays the consequences of your configuration. You should use this feature cautiously. Using the old master MAC address elsewhere in the domain could result in lost traffic.

You can set the time period from 0 to 60 minutes.

- If you enter the command with no value, the default delay is 4 minutes. We recommend that you always enter a value. The time delay appears in the configuration file with an explicit timer value of 4 minutes.
- If you enter **0**, the stack MAC address of the previous master is used until you enter the **no stack-mac persistent timer** global configuration command, which changes the stack MAC address to that of the current master. If you do not enter this command, the stack MAC address does not change.
- If you enter a time delay of 1 to 60 minutes, the stack MAC address of the previous master is used until the configured time period expires or until you enter the **no stack-mac persistent timer** command.


If the previous master does not rejoin the stack during this period, the stack uses the MAC address of the new master as the stack MAC address.



Note

If the entire switch stack reloads, it acquires the MAC address of the master as the stack MAC address.

Beginning in privileged EXEC mode, follow these steps to enable persistent MAC address. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	stack-mac persistent timer [0 <i>time-value</i>]	<p>Enable a time delay after a stack-master change before the stack MAC address changes to that of the new stack master. If the previous stack master rejoins the stack during this period, the stack uses that MAC address as the stack MAC address.</p> <ul style="list-style-type: none"> Enter the command with no value to set the default delay of 4 minutes. We recommend that you always configure a value. Enter 0 to use the MAC address of the current master indefinitely. Enter a <i>time-value</i> from 1 to 60 to configure the time period (in minutes) before the stack MAC address changes to the new master. <p> Caution When you enter this command, a warning states that traffic might be lost if the old master MAC address appears elsewhere in the network domain.</p> <p>If you enter the no stack-mac persistent timer command after a new stack master takes over, before the time expires, the stack uses the current master MAC address.</p>
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config or show switch	<p>Verify that the stack MAC address timer is enabled.</p> <p>The output shows <code>stack-mac persistent timer</code> and the time in minutes.</p> <p>The output shows <code>Mac persistency wait time</code> with the number of minutes configured and the stack MAC address.</p>
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no stack-mac persistent timer** global configuration command to disable the persistent MAC address feature.

This example shows how to configure the persistent MAC address feature for a 7-minute time delay and to verify the configuration:

```
Switch(config)# stack-mac persistent timer 7
WARNING: The stack continues to use the base MAC of the old Master
WARNING: as the stack MAC after a master switchover until the MAC
WARNING: persistency timer expires. During this time the Network
WARNING: Administrators must make sure that the old stack-mac does
WARNING: not appear elsewhere in this network domain. If it does,
WARNING: user traffic may be blackholed.
Switch(config)# end
```

```

Switch# show switch
Switch/Stack Mac Address : 0016.4727.a900
Mac persistency wait time: 7 mins

                H/W   Current
Switch#  Role   Mac Address      Priority Version  State
-----
*1       Master  0016.4727.a900   1         0         Ready

```

Assigning Stack Member Information

- [Assigning a Member Number, page 1-19](#) (optional)
- [Setting the Member Priority Value, page 1-20](#) (optional)
- [Provisioning a New Member for a Stack, page 1-20](#) (optional)

Assigning a Member Number



Note This task is available only from the master.

Beginning in privileged EXEC mode, follow these steps to assign a member number to a member. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	switch <i>current-stack-member-number</i> renumber <i>new-stack-member-number</i>	Specify the current member number and the new member number for the member. The range is 1 to 4. You can display the current member number by using the show switch user EXEC command.
Step 3	end	Return to privileged EXEC mode.
Step 4	reload slot <i>stack-member-number</i>	Reset the stack member.
Step 5	show switch	Verify the stack member number.
Step 6	copy running-config startup-config	Save your entries in the configuration file.

Setting the Member Priority Value



Note This task is available only from the master.

Beginning in privileged EXEC mode, follow these steps to assign a priority value to a member: This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	switch <i>stack-member-number</i> priority <i>new-priority-number</i>	Specify the member number and the new priority for the member. The member number range is 1 to 4. The priority value range is 1 to 15. You can display the current priority value by using the show switch user EXEC command. The new priority value takes effect immediately but does not affect the current master until the current master or the stack resets.
Step 3	end	Return to privileged EXEC mode.
Step 4	reload slot <i>stack-member-number</i>	Reset the member, and apply this configuration.
Step 5	show switch <i>stack-member-number</i>	Verify the member priority value.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

You can also set the SWITCH_PRIORITY environment variable. For more information, see the [“Controlling Environment Variables” section on page 1-20](#).

Provisioning a New Member for a Stack



Note This task is available only from the master.

Beginning in privileged EXEC mode, follow these steps to provision a new member for a stack. This procedure is optional.

	Command	Purpose
Step 1	show switch	Display summary information about the stack.
Step 2	configure terminal	Enter global configuration mode.
Step 3	switch <i>stack-member-number</i> provision <i>type</i>	Specify the member number for the provisioned switch. By default, no switches are provisioned. For <i>stack-member-number</i> , the range is 1 to 4. Enter a member number that is not already used in the stack. See Step 1. For <i>type</i> , enter the model number of the member.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify the correct numbering of interfaces in the configuration.

	Command	Purpose
Step 6	<code>show switch stack-member-number</code>	Verify the status of the provisioned switch. For <i>stack-member-number</i> , enter the same number as in Step 2.
Step 7	<code>copy running-config startup-config</code>	(Optional) Save your entries in the configuration file.

To remove provisioned information and to avoid receiving an error message, remove the specified switch from the stack before you use the **no** form of this command.

This example shows how to provision a switch with a stack member number of 2 for the stack. The **show running-config** command output shows the interfaces associated with the provisioned switch:

```
Switch(config)# switch 2 provision
Switch(config)# end
Switch# show running-config | include switch 2
!
interface GigabitEthernet2/0/1
!
interface GigabitEthernet2/0/2
!
interface GigabitEthernet2/0/3
<output truncated>
```

Changing the Stack Membership

If you remove powered-on members but do not want to partition the stack:

-
- Step 1 Power off the newly created stacks.
 - Step 2 Reconnect them to the original stack through their stack ports.
 - Step 3 Power on the switches.
-

Accessing the CLI of a Specific Member



Note

This task is only for debugging purposes, and is only available from the master.

You can access all or specific members by using the **remote command** `{all | stack-member-number}` privileged EXEC command. The stack member number range is 1 to 4.

You can access specific members by using the **session** `stack-member-number` privileged EXEC command. The member number is appended to the system prompt. For example, the prompt for member 2 is `Switch-2#`, and system prompt for the master is `Switch#`. Enter `exit` to return to the CLI session on the master. Only the **show** and **debug** commands are available on a specific member.

For more information, see the [“Using Interface Configuration Mode”](#) section on page 1-15.

Displaying Stack Information

To display saved configuration changes after resetting a specific member or the stack, use these privileged EXEC commands:

Table 1-4 Commands for Displaying Stack Information

Command	Description
show controller ethernet-controller stack port [1 2]	Display stack port counters (or per-interface and per-stack port send and receive statistics read from the hardware).
show platform stack passive-links all	Display all stack information, such as the stack protocol version.
show switch	Display summary information about the stack, including the status of provisioned switches and switches in version-mismatch mode.
show switch <i>stack-member-number</i>	Display information about a specific member.
show switch detail	Display detailed information about the stack ring.
show switch neighbors	Display the stack neighbors.
show switch stack-ports	Display port information for the stack.

Troubleshooting Stacks

- [Manually Disabling a Stack Port, page 1-22](#)
- [Re-Enabling a Stack Port While Another Member Starts, page 1-23](#)
- [Understanding the show switch stack-ports summary Output, page 1-23](#)

Manually Disabling a Stack Port

If a stack port is flapping and causing instability in the stack ring, to disable the port, enter the **switch stack-member-number stack port port-number disable** privileged EXEC command. To re-enable the port, enter the **switch stack-member-number stack port port-number enable** command.



Note

Be careful when using the **switch stack-member-number stack port port-number disable** command. When you disable the stack port, .

- A stack is in the *full-ring* state when all members are connected through the stack ports and are in the ready state.
- The stack is in the *partial-ring* state when
 - All members are connected through the stack ports, but some all are not in the ready state.
 - Some members are not connected through the stack ports.

When you enter the **switch stack-member-number stack port port-number disable** privileged EXEC command and

- The stack is in the full-ring state, you can disable only one stack port. This message appears:
Enabling/disabling a stack port may cause undesired stack changes. Continue?[confirm]

- The stack is in the partial-ring state, you cannot disable the port. This message appears:
Disabling stack port not allowed with current stack configuration.

Re-Enabling a Stack Port While Another Member Starts

Stack Port 1 on Switch 1 is connected to Port 2 on Switch 4. If Port 1 is flapping, disable Port 1 with the **switch 1 stack port 1 disable** privileged EXEC command.

While Port 1 on Switch 1 is disabled and Switch 1 is still powered on:

1. Disconnect the stack cable between Port 1 on Switch 1 and Port 2 on Switch 4.
2. Remove Switch 4 from the stack.
3. Add a switch to replace Switch 4 and assign it switch-number 4.
4. Reconnect the cable between Port 1 on Switch 1 and Port 2 on Switch 4 (the replacement switch).
5. Re-enable the link between the switches. Enter the **switch 1 stack port 1 enable** privileged EXEC command to enable Port 1 on Switch 1.
6. Power on Switch 4.



Caution

Powering on Switch 4 before enabling the Port 1 on Switch 1 might cause one of the switches to reload.

If Switch 4 is powered on first, you might need to enter the **switch 1 stack port 1 enable** and the **switch 4 stack port 2 enable** privileged EXEC commands to bring up the link.

Understanding the show switch stack-ports summary Output

Only Port 1 on stack member 2 is disabled.

```
Switch# show switch stack-ports summary
Switch#/  Stack  Neighbor  Cable  Link  Link  Sync  #  In
Port#    Port           Length  OK    Active  OK    Changes  Loopback
        Status
-----  -
1/1      OK           3       50 cm  Yes   Yes   Yes    1    No
1/2      Down        None    3 m    Yes   No    Yes    1    No
2/1      Down        None    3 m    Yes   No    Yes    1    No
2/2      OK           3       50 cm  Yes   Yes   Yes    1    No
3/1      OK           2       50 cm  Yes   Yes   Yes    1    No
3/2      OK           1       50 cm  Yes   Yes   Yes    1    No
```

Table 1-5 show switch stack-ports summary Command Output

Field	Description
Switch#/Port#	Member number and its stack port number.
Stack Port Status	<ul style="list-style-type: none"> • Absent—No cable is detected on the stack port. • Down—A cable is detected, but either no connected neighbor is up, or the stack port is disabled. • OK—A cable is detected, and the connected neighbor is up.
Neighbor	Switch number of the active member at the other end of the stack cable.

Table 1-5 *show switch stack-ports summary Command Output (continued)*

Field	Description
Cable Length	Valid lengths are 50 cm, 1 m, or 3 m. If the switch cannot detect the cable length, the value is <i>no cable</i> . The cable might not be connected, or the link might be unreliable.
Link OK	This shows if the link is stable. The <i>link partner</i> is a stack port on a neighbor switch. <ul style="list-style-type: none"> No—The link partner receives invalid protocol messages from the port. Yes—The link partner receives valid protocol messages from the port.
Link Active	This shows if the stack port is in the same state as its link partner. <ul style="list-style-type: none"> No—The port cannot send traffic to the link partner. Yes—The port can send traffic to the link partner.
Sync OK	<ul style="list-style-type: none"> No—The link partner does not send valid protocol messages to the stack port. Yes—The link partner sends valid protocol messages to the port.
# Changes to LinkOK	This shows the relative stability of the link. If a large number of changes occur in a short period of time, link flapping can occur.
In Loopback	<ul style="list-style-type: none"> No—At least one stack port on the member has an attached stack cable. Yes—None of the stack ports on the member has an attached stack cable.



CHAPTER 1

Configuring SDM Templates

The Catalyst 2960 and 2960-S switch command reference has command syntax and usage information. Unless otherwise noted, the term *switch* refers to a standalone switch and a switch stack.



Note

Configuring an SDM template is not required on Catalyst 2960-S switches unless you configure static routing on the switch. A Catalyst 2960-S switch running the LAN base image uses a default template that includes maximum resources for all supported features.

- [Understanding the SDM Templates, page 1-1](#)
- [Configuring the Switch SDM Template, page 1-3](#)
- [.Displaying the SDM Templates, page 1-4](#)

Understanding the SDM Templates



Note

The SDM template used by Catalyst 2960-S running LAN Lite image is a default templates and is not configurable. Catalyst 2960-S switches running the LAN base image support only a default template and the lanbase-routing template.

You can use SDM templates to configure system resources in the Catalyst 2960 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 use the default template to balance resources.

To allocate ternary content addressable memory (TCAM) resources for different usages, the switch SDM templates prioritize system resources to optimize support for certain features. You can select SDM templates to optimize these features on the Catalyst 2960 switch:

- Default—The default template gives balance to all functions.
- Dual—The dual IPv4 and IPv6 template allows the switch to be used in dual stack environments (supporting both IPv4 and IPv6). Using the dual stack templates results in less TCAM capacity allowed for each resource. Do not use them if you plan to forward only IPv4 traffic.



Note

The dual IPv4 and IPv6 template is not supported on Catalyst 2960 switches running the LAN Lite image or on Catalyst 2960-S switches.

- LAN base routing—The lanbase-routing template supports IPv4 unicast routes for configuring static routing SVIs



Note The lanbase-routing template is supported only on switches running Cisco IOS Release 12.2(55)SE or later and the LAN base image.

- QoS—The QoS template maximizes system resources for quality of service (QoS) access control entries (ACEs).

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

Resource	Default Catalyst 2960	Default Catalyst 2960-S	QoS Catalyst 2960 Only	Dual Catalyst 2960 Only	LAN base routing
Unicast MAC addresses	8 K	8K	8 K	8 K	4K
IPv4 IGMP groups	256	256	256	256	256
IPv4 unicast routes	0	256	0	0	4.25K
• Directly connected hosts	0		0	0	4K
• Indirect routes	0		0	0	256
IPv6 multicast groups	0		0	0	0
Directly connected IPv6 addresses	0		0	0	0
Indirect IPv6 unicast routes	0		0	0	0
IPv4 policy-based routing aces	0		0	0	0
IPv4 MAC QoS ACEs	128	384	384	0	128
IPv4 MAC security ACEs	384	384	128	256	384
IPv6 policy-based routing aces	0		0	0	0
IPv6 QoS ACEs	0		0	0	0
IPv6 security ACEs	0	128	0	0	0

The rows in the tables 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.

SDM Templates and Switch Stacks



Note Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

All stack members use the same SDM template that is stored on the stack master. When a new switch is added to a stack, as with the switch configuration and VLAN database files, the SDM configuration that is stored on the stack master overrides the template configured on an individual switch. For more information about stacking, see [Chapter 1, “Managing Switch Stacks.”](#)

You can use the **show switch** privileged EXEC command to see if any stack members are in SDM mismatch mode. This example shows the output from the **show switch** privileged EXEC command when an SDM mismatch exists:

```
Switch# show switch

Switch# Role      Mac Address      Priority    Current
-----
*2      Master    000a.fdfd.0100   5          Ready
4       Member    0003.fd63.9c00   5          SDM Mismatch
```

This is an example of a syslog message notifying the stack master that a stack member is in SDM mismatch mode:

```
!!!!!!! SDM MISMATCH !!!!!!!
Master Template is lanbase-routing & Local Template is default
Reloading because of sdm template mismatch
Please reboot the switch
```

Configuring the Switch SDM Template

- [Default SDM Template, page 1-3](#)
- [SDM Template Configuration Guidelines, page 1-3](#)
- [Setting the SDM Template, page 1-4](#)

Default SDM Template

The default template for the Catalyst 2960 and 2960-S switches is the default desktop template.

SDM Template Configuration Guidelines



Note

You configure SDM template on Catalyst 2960 switches. A Catalyst 2960-S switch running the LAN base image uses the desktop default template that includes maximum resources for all supported features. However, to enable static routing, you must configure the lanbase-routing template.

- When you select and configure SDM templates, you must reload the switch for the configuration to take effect.
- Do not use the routing template if you do not have routing enabled on your switch. The **sdm prefer lanbase routing** global configuration command prevents other features from using the memory allocated to unicast routing in the routing template.
- If you try to configure IPv6 features without first selecting a dual IPv4 and IPv6 template, a warning message appears.



Note

The dual template is not supported on switches running the LAN lite image and is not required on Catalyst 2960-S switches.

- Using the dual stack templates results in less TCAM capacity allowed for each resource, so do not use it if you plan to forward only IPv4 traffic.

Setting the SDM Template

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

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	sdm prefer { default dual-ipv4-and-ipv6 default lanbase-routing qos }	Specify the SDM template to be used on the switch: The keywords have these meanings: <ul style="list-style-type: none"> • default—Gives balance to all functions. • dual-ipv4-and-ipv6 default—Allows the switch to be used in dual stack environments (supporting both IPv4 and IPv6). • lanbase-routing—Supports configuring unicast routes for static routing on SVIs. • qos—Maximizes system resources for QoS ACEs. Use the no sdm prefer command to set the switch to the default template. The default template balances the use of system resources. Note The Catalyst 2960-S switch supports only the default and lanbase-routing templates.
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 in use and the template that becomes active after a reload.

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 | dual-ipv4-and-ipv6 default | lanbase-routing | qos]** privileged EXEC command to display the resource numbers supported by the specified template.



Note

The Catalyst 2960-S switch supports only the default and lanbase-routing templates.

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 "lanbase-routing" template.
The selected template optimizes the resources in
the switch to support this level of features for
8 routed interfaces and 255 VLANs.

number of unicast mac addresses:          4K
number of IPv4 IGMP groups + multicast routes: 0.25K
number of IPv4 unicast routes:           0.75K
  number of directly-connected IPv4 hosts: 0.75K
  number of indirect IPv4 routes:         1K
number of IPv4 policy based routing aces: 0
number of IPv4/MAC qos aces:             0.125k
number of IPv4/MAC security aces:        0.375k
```




CHAPTER 1

Configuring Switch-Based Authentication

This chapter describes how to configure switch-based authentication on the Catalyst 2960 and 2960-S switches. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

This chapter consists of these sections:

- [Preventing Unauthorized Access to Your Switch, page 1-1](#)
- [Protecting Access to Privileged EXEC Commands, page 1-2](#)
- [Controlling Switch Access with TACACS+, page 1-10](#)
- [Controlling Switch Access with RADIUS, page 1-17](#)
- [Configuring the Switch for Local Authentication and Authorization, page 1-40](#)
- [Configuring the Switch for Secure Shell, page 1-41](#)
- [Configuring the Switch for Secure Socket Layer HTTP, page 1-46](#)
- [Configuring the Switch for Secure Copy Protocol, page 1-52](#)

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 1-2](#).
- 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 1-6.

- 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 1-10.

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, see the *Cisco IOS Security Command Reference, Release 12.4* on Cisco.com.

These sections contain this configuration information:

- [Default Password and Privilege Level Configuration, page 1-2](#)
- [Setting or Changing a Static Enable Password, page 1-3](#)
- [Protecting Enable and Enable Secret Passwords with Encryption, page 1-3](#)
- [Disabling Password Recovery, page 1-5](#)
- [Setting a Telnet Password for a Terminal Line, page 1-6](#)
- [Configuring Username and Password Pairs, page 1-6](#)
- [Configuring Multiple Privilege Levels, page 1-7](#)

Default Password and Privilege Level Configuration

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

Table 1-1 *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 *11u2c3k4y5*. The password is not encrypted and provides access to level 15 (traditional privileged EXEC mode access):

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

Protecting Enable and Enable Secret Passwords with Encryption

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

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

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

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

	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. <ul style="list-style-type: none"> • (Optional) For <i>level</i>, the range is from 0 to 15. Level 1 is normal user EXEC mode privileges. The default level is 15 (privileged EXEC mode privileges). • For <i>password</i>, specify a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. By default, no password is defined. • (Optional) For <i>encryption-type</i>, only type 5, a Cisco proprietary encryption algorithm, is available. If you specify an encryption type, you must provide an encrypted password—an encrypted password that you copy from another switch configuration. <p>Note If you specify an encryption type and then enter a clear text password, you can not re-enter privileged EXEC mode. You cannot recover a lost encrypted password by any method.</p>
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 1-7](#).

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 `1FaD0$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 1-3.

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 command output.

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



Note

Disabling password recovery will not work if you have set the switch to boot up 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. <ul style="list-style-type: none"> For <i>name</i>, specify the user ID as one word. Spaces and quotation marks are not allowed. (Optional) For <i>level</i>, specify the privilege level the user has after gaining access. The range is 0 to 15. Level 15 gives privileged EXEC mode access. Level 1 gives user EXEC mode access. For <i>encryption-type</i>, enter 0 to specify that an unencrypted password will follow. Enter 7 to specify that a hidden password will follow. For <i>password</i>, specify the password the user must enter to gain access to the switch. The password must be from 1 to 25 characters, can contain embedded spaces, and must be the last option specified in the username command.
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 Cisco IOS 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.

These sections contain this configuration information:

- [Setting the Privilege Level for a Command, page 1-8](#)
- [Changing the Default Privilege Level for Lines, page 1-9](#)
- [Logging into and Exiting a Privilege Level, page 1-9](#)

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

Controlling Switch Access with TACACS+

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

Beginning with Cisco IOS Release 12.2(58)SE, the switch supports TACACS+ for IPv6. Information is in the “[TACACS+ Over an IPv6 Transport](#)” section of the “Implementing ADSL for IPv6” chapter in the *Cisco IOS XE IPv6 Configuration Guide, Release 2*.

For information about configuring this feature, see the “[Configuring TACACS+ over IPv6](#)” section of the “Implementing ADSL for IPv6” chapter in the *Cisco IOS XE IPv6 Configuration Guide, Release 2*.



Note

For complete syntax and usage information for the commands used in this section, see the *Cisco IOS Security Command Reference, Release 12.4* and the *Cisco IOS IPv6 Command Reference*.



Note

For complete syntax and usage information for the commands used in this section, see the *Cisco IOS Security Command Reference, Release 12.4*.

These sections contain this configuration information:

- [Understanding TACACS+, page 1-10](#)
- [TACACS+ Operation, page 1-12](#)
- [Configuring TACACS+, page 1-12](#)
- [Displaying the TACACS+ Configuration, page 1-17](#)

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.



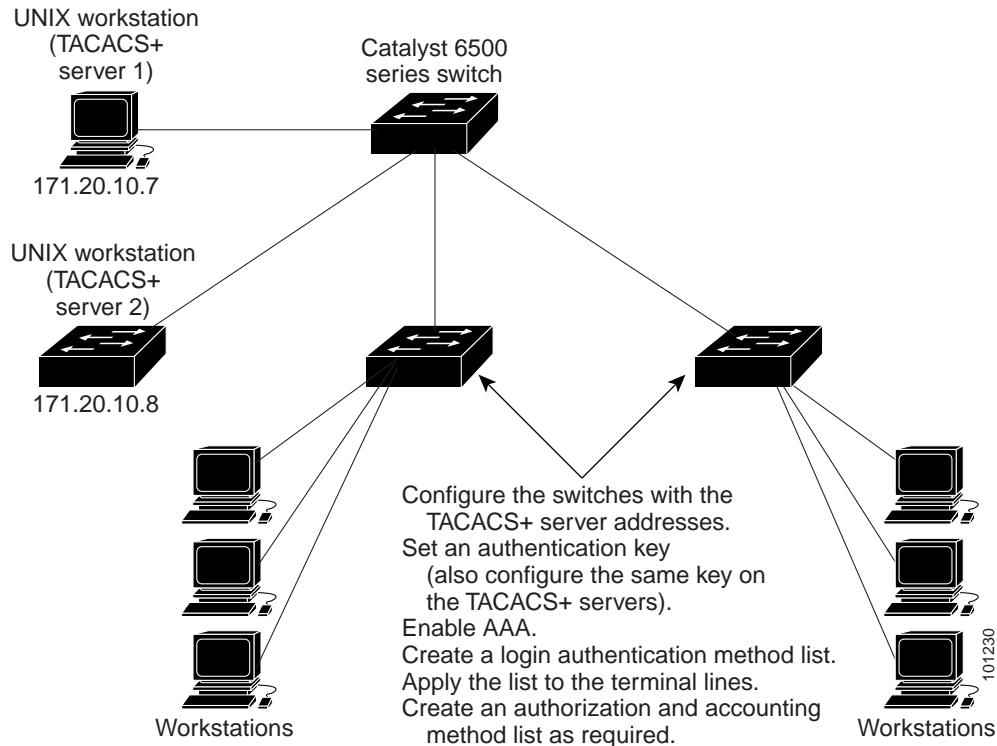
Note

We recommend a redundant connection between a switch stack and the TACACS+ server. This is to help ensure that the TACACS+ server remains accessible in case one of the connected stack members is removed from the switch stack.

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 1-1](#).

Figure 1-1 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 to show 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 dialog 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.

These sections contain this configuration information:

- [Default TACACS+ Configuration, page 1-13](#)
- [Identifying the TACACS+ Server Host and Setting the Authentication Key, page 1-13](#)
- [Configuring TACACS+ Login Authentication, page 1-14](#)
- [Configuring TACACS+ Authorization for Privileged EXEC Access and Network Services, page 1-16](#)
- [Starting TACACS+ Accounting, page 1-17](#)

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. <ul style="list-style-type: none"> • For <i>hostname</i>, specify the name or IP address of the host. • (Optional) For port <i>integer</i>, specify a server port number. The default is port 49. The range is 1 to 65535. • (Optional) For timeout <i>integer</i>, specify a time in seconds the switch waits for a response from the daemon before it times out and declares an error. The default is 5 seconds. The range is 1 to 1000 seconds. • (Optional) For key <i>string</i>, specify the encryption key for encrypting and decrypting all traffic between the switch and the TACACS+ daemon. You must configure the same key on the TACACS+ daemon for encryption to be successful.
Step 3	aaa new-model	Enable AAA.

	Command	Purpose
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.
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>]	<p>Create a login authentication method list.</p> <ul style="list-style-type: none"> To create a default list that is used when a named list is <i>not</i> specified in the login authentication command, use the default keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all 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. <p>Select one of these methods:</p> <ul style="list-style-type: none"> enable—Use the enable password for authentication. Before you can use this authentication method, you must define an enable password by using the enable password global configuration command. group tacacs+—Uses TACACS+ authentication. Before you can use this authentication method, you must configure the TACACS+ server. For more information, see the “Identifying the TACACS+ Server Host and Setting the Authentication Key” section on page 1-13. 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> }	<p>Apply the authentication list to a line or set of lines.</p> <ul style="list-style-type: none"> If you specify default, use the default list created with the aaa authentication login command. For <i>list-name</i>, specify the list created with the aaa authentication login command.
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.

**Note**

To secure the switch for HTTP access by using AAA methods, you must configure the switch with the **ip http authentication aaa** global configuration command. Configuring AAA authentication does not secure the switch for HTTP access by using AAA methods.

For more information about the **ip http authentication** command, see the *Cisco IOS Security Command Reference, Release 12.4* on Cisco.com.

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.

Establishing a Session with a Router if the AAA Server is Unreachable



Note

To configure this command, the switch must be running the LAN Base image.

The **aaa accounting system guarantee-first** command guarantees system accounting as the first record, which is the default condition. In some situations, users might be prevented from starting a session on the console or terminal connection until after the system reloads, which can take more than 3 minutes.

To establish a console or Telnet session with the router if the AAA server is unreachable when the router reloads, use the **no aaa accounting system guarantee-first** 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.

Beginning with Cisco IOS Release 12.2(58)SE, the switch supports RADIUS for IPv6. Information is in the “RADIUS Over IPv6” section of the “Implementing ADSL for IPv6” chapter in the *Cisco IOS XE IPv6 Configuration Guide, Release 2*. For information about configuring this feature, see the “Configuring the NAS” section in the “Implementing ADSL for IPv6” chapter in the *Cisco IOS XE IPv6 Configuration Guide, Release 2*.

**Note**

For complete syntax and usage information for the commands used in this section, see the *Cisco IOS Security Command Reference* and the *Cisco IOS IPv6 Command Reference*.

These sections contain this configuration information:

- [Understanding RADIUS, page 1-18](#)
- [RADIUS Operation, page 1-19](#)
- [RADIUS Change of Authorization, page 1-20](#)
- [Configuring RADIUS, page 1-26](#)
- [Displaying the RADIUS Configuration, page 1-40](#)

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, see the RADIUS server documentation.

**Note**

We recommend a redundant connection between a switch stack and the RADIUS server. This is to help ensure that the RADIUS server remains accessible in case one of the connected stack members is removed from the switch stack.

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 1-2 on page 1-19](#).
- 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 1, “Configuring IEEE 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 1-2 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

RADIUS Change of Authorization

To use this feature, the switch must be running the LAN Base or IP Base image.

This section provides an overview of the RADIUS interface including available primitives and how they are used during a Change of Authorization (CoA).

- [Overview, page 1-20](#)
- [Change-of-Authorization Requests, page 1-21](#)
- [CoA Request Response Code, page 1-22](#)
- [CoA Request Commands, page 1-23](#)
- [Session Reauthentication, page 1-24](#)
- [Stacking Guidelines for Session Termination, page 1-25](#)

Overview

A standard RADIUS interface is typically used in a pulled model where the request originates from a network attached device and the response come from the queried servers. Catalyst switches support the RADIUS Change of Authorization (CoA) extensions defined in RFC 5176 that are typically used in a pushed model and allow for the dynamic reconfiguring of sessions from external authentication, authorization, and accounting (AAA) or policy servers.

Beginning with Cisco IOS Release 12.2(52)SE, the switch supports these per-session CoA requests:

- Session reauthentication
- Session termination
- Session termination with port shutdown
- Session termination with port bounce

This feature is integrated with the Cisco Secure Access Control Server (ACS) 5.1. For information about ACS:

http://www.cisco.com/en/US/products/ps9911/tsd_products_support_series_home.html

The RADIUS interface is enabled by default on Catalyst switches. However, some basic configuration is required for these attributes:

- Security and Password—See the [“Preventing Unauthorized Access to Your Switch”](#) section in the “Configuring Switch-Based Authentication” chapter in the *Catalyst 3750 Switch Software Configuration Guide, Cisco Release 12.2(50)SE*.
- Accounting—See the [“Starting RADIUS Accounting”](#) section in the “Configuring Switch-Based Authentication” chapter in the *Catalyst 3750 Switch Software Configuration Guide, 12.2(50)SE*.

Change-of-Authorization Requests

Change of Authorization (CoA) requests, as described in RFC 5176, are used in a push model to allow for session identification, host reauthentication, and session termination. The model is comprised of one request (CoA-Request) and two possible response codes:

- CoA acknowledgement (ACK) [CoA-ACK]
- CoA non-acknowledgement (NAK) [CoA-NAK]

The request is initiated from a CoA client (typically a RADIUS or policy server) and directed to the switch that acts as a listener.

This section includes these topics:

- [CoA Request Response Code](#)
- [CoA Request Commands](#)
- [Session Reauthentication](#)

RFC 5176 Compliance

The Disconnect Request message, which is also referred to as Packet of Disconnect (POD), is supported by the switch for session termination.

[Table 1-2](#) shows the IETF attributes are supported for this feature.

Table 1-2 Supported IETF Attributes

Attribute Number	Attribute Name
24	State
31	Calling-Station-ID
44	Acct-Session-ID
80	Message-Authenticator
101	Error-Cause

[Table 1-3](#) shows the possible values for the Error-Cause attribute.

Table 1-3 Error-Cause Values

Value	Explanation
201	Residual Session Context Removed
202	Invalid EAP Packet (Ignored)
401	Unsupported Attribute
402	Missing Attribute
403	NAS Identification Mismatch
404	Invalid Request
405	Unsupported Service
406	Unsupported Extension
407	Invalid Attribute Value

Table 1-3 Error-Cause Values (continued)

Value	Explanation
501	Administratively Prohibited
502	Request Not Routable (Proxy)
503	Session Context Not Found
504	Session Context Not Removable
505	Other Proxy Processing Error
506	Resources Unavailable
507	Request Initiated
508	Multiple Session Selection Unsupported

Preconditions

To use the CoA interface, a session must already exist on the switch. CoA can be used to identify a session and enforce a disconnect request. The update affects only the specified session.

CoA Request Response Code

The CoA Request response code can be used to convey a command to the switch. The supported commands are listed in [Table 1-4 on page 1-23](#).

Session Identification

For disconnect and CoA requests targeted at a particular session, the switch locates the session based on one or more of the following attributes:

- Calling-Station-Id (IETF attribute 31 which contains the host MAC address)
- Audit-Session-Id (Cisco VSA)
- Acct-Session-Id (IETF attribute 44)


Unless all session identification attributes included in the CoA message match the session, the switch returns a Disconnect-NAK or CoA-NAK with the Invalid Attribute Value error-code attribute.

For disconnect and CoA requests targeted to a particular session, any one of these session identifiers can be used:

- Calling-Station-ID (IETF attribute 31, which should contain the MAC address)
- Audit-Session-ID (Cisco vendor-specific attribute)
- Accounting-Session-ID (IETF attribute 44).

If more than one session identification attribute is included in the message, all the attributes must match the session or the switch returns a Disconnect- negative acknowledgement (NAK) or CoA-NAK with the error code *Invalid Attribute Value*.

The packet format for a CoA Request code as defined in RFC 5176 consists of the fields: Code, Identifier, Length, Authenticator, and Attributes in Type:Length:Value (TLV) format.



The attributes field is used to carry Cisco VSAs.

CoA ACK Response Code

If the authorization state is changed successfully, a positive acknowledgement (ACK) is sent. The attributes returned within CoA ACK will vary based on the CoA Request and are discussed in individual CoA Commands.

CoA NAK Response Code

A negative acknowledgement (NAK) indicates a failure to change the authorization state and can include attributes that indicate the reason for the failure. Use **show** commands to verify a successful CoA.

CoA Request Commands

This section includes:

- [Session Reauthentication](#)
- [Session Termination](#)
- [CoA Disconnect-Request](#)
- [CoA Request: Disable Host Port](#)
- [CoA Request: Bounce-Port](#)

Beginning with Cisco IOS Release 12.2(52)SE, the switch supports the commands shown in [Table 1-4](#).

Table 1-4 CoA Commands Supported on the Switch

Command ¹	Cisco VSA
Reauthenticate host	Cisco:Avpair="subscriber:command=reauthenticate"
Terminate session	This is a standard disconnect request that does not require a VSA.
Bounce host port	Cisco:Avpair="subscriber:command=bounce-host-port"
Disable host port	Cisco:Avpair="subscriber:command=disable-host-port"

1. All CoA commands must include the session identifier between the switch and the CoA client.

Session Reauthentication

The AAA server typically generates a session reauthentication request when a host with an unknown identity or posture joins the network and is associated with a restricted access authorization profile (such as a guest VLAN). A reauthentication request allows the host to be placed in the appropriate authorization group when its credentials are known.

To initiate session authentication, the AAA server sends a standard CoA-Request message which contains a Cisco vendor-specific attribute (VSA) in this form:

Cisco:Avpair="subscriber:command=reauthenticate" and one or more session identification attributes.

The current session state determines the switch response to the message. If the session is currently authenticated by IEEE 802.1x, the switch responds by sending an Extensible Authentication Protocol over LAN (EAPoL) RequestId message to the server.

If the session is currently authenticated by MAC authentication bypass (MAB), the switch sends an access-request to the server, passing the same identity attributes used for the initial successful authentication.

If session authentication is in progress when the switch receives the command, the switch terminates the process, and restarts the authentication sequence, starting with the method configured to be attempted first.

If the session is not yet authorized, or is authorized via guest VLAN, or critical VLAN, or similar policies, the reauthentication message restarts the access control methods, beginning with the method configured to be attempted first. The current authorization of the session is maintained until the reauthentication leads to a different authorization result.

Session Reauthentication in a Switch Stack

When a switch stack receives a session reauthentication message:

- It checkpoints the need for a re-authentication before returning an acknowledgement (ACK).
- It initiates reauthentication for the appropriate session.
- If authentication completes with either success or failure, the signal that triggered the reauthentication is removed from the stack member.
- If the stack master fails before authentication completes, reauthentication is initiated after stack master switch-over based on the original command (which is subsequently removed).
- If the stack master fails before sending an ACK, the new stack master treats the re-transmitted command as a new command.

Session Termination

There are three types of CoA requests that can trigger session termination. A CoA Disconnect-Request terminates the session, without disabling the host port. This command causes re-initialization of the authenticator state machine for the specified host, but does not restrict that host's access to the network.

To restrict a host's access to the network, use a CoA Request with the *Cisco:Avpair="subscriber:command=disable-host-port"* VSA. This command is useful when a host is known to be causing problems on the network, and you need to immediately block network access for the host. When you want to restore network access on the port, re-enable it using a non-RADIUS mechanism.

When a device with no supplicant, such as a printer, needs to acquire a new IP address (for example, after a VLAN change), terminate the session on the host port with port-bounce (temporarily disable and then re-enable the port).

CoA Disconnect-Request

This command is a standard Disconnect-Request. Because this command is session-oriented, it must be accompanied by one or more of the session identification attributes described in the [“Session Identification” section on page 1-22](#). If the session cannot be located, the switch returns a Disconnect-NAK message with the “Session Context Not Found” error-code attribute. If the session is located, the switch terminates the session. After the session has been completely removed, the switch returns a Disconnect-ACK.

If the switch fails-over to a standby switch before returning a Disconnect-ACK to the client, the process is repeated on the new active switch when the request is re-sent from the client. If the session is not found following re-sending, a Disconnect-ACK is sent with the “Session Context Not Found” error-code attribute.

CoA Request: Disable Host Port

This command is carried in a standard CoA-Request message that has this new VSA:

```
Cisco:Avpair="subscriber:command=disable-host-port"
```

Because this command is session-oriented, it must be accompanied by one or more of the session identification attributes described in the [“Session Identification” section on page 1-22](#). If the session cannot be located, the switch returns a CoA-NAK message with the “Session Context Not Found” error-code attribute. If the session is located, the switch disables the hosting port and returns a CoA-ACK message.

If the switch fails before returning a CoA-ACK to the client, the process is repeated on the new active switch when the request is re-sent from the client. If the switch fails after returning a CoA-ACK message to the client but before the operation has completed, the operation is restarted on the new active switch.



Note

A Disconnect-Request failure following command re-sending could be the result of either a successful session termination before change-over (if the Disconnect-ACK was not sent) or a session termination by other means (for example, a link failure) that occurred after the original command was issued and before the standby switch became active.

CoA Request: Bounce-Port

This command is carried in a standard CoA-Request message that contains this VSA:

```
Cisco:Avpair="subscriber:command=bounce-host-port"
```

Because this command is session-oriented, it must be accompanied by one or more of the session identification attributes described in the [“Session Identification” section on page 1-22](#). If the session cannot be located, the switch returns a CoA-NAK message with the “Session Context Not Found” error-code attribute. If the session is located, the switch disables the hosting port for a period of 10 seconds, re-enables it (port-bounce), and returns a CoA-ACK.

If the switch fails before returning a CoA-ACK to the client, the process is repeated on the new active switch when the request is re-sent from the client. If the switch fails after returning a CoA-ACK message to the client but before the operation has completed, the operation is re-started on the new active switch.

Stacking Guidelines for Session Termination

No special handling is required for CoA Disconnect-Request messages in a switch stack.

Stacking Guidelines for CoA-Request Bounce-Port

Because the **bounce-port** command is targeted at a session, not a port, if the session is not found, the command cannot be executed.

When the Auth Manager command handler on the stack master receives a valid **bounce-port** command, it checkpoints this information before returning a CoA-ACK message:

- Need for a port-bounce
- Port-ID (found in the local session context)

The switch initiates a port-bounce (disables the port for 10 seconds, then re-enables it).

If the port-bounce is successful, the signal that triggered the port-bounce is removed from the standby stack master.

If the stack master fails before the port-bounce completes, a port-bounce is initiated after stack master change-over based on the original command (which is subsequently removed).

If the stack master fails before sending a CoA-ACK message, the new stack master treats the re-sent command as a new command.

Stacking Guidelines for CoA-Request Disable-Port

Because the **disable-port** command is targeted at a session, not a port, if the session is not found, the command cannot be executed.

When the Auth Manager command handler on the stack master receives a valid **disable-port** command, it verifies this information before returning a CoA-ACK message:

- Need for a port-disable
- Port-ID (in the local session context)

The switch attempts to disable the port.

If the port-disable operation is successful, the signal that triggered the port-disable is removed from the standby stack master.

If the stack master fails before the port-disable operation completes, the port is disabled after stack master change-over based on the original command (which is subsequently removed).

If the stack master fails before sending a CoA-ACK message, the new stack master treats the re-sent command as a new command.

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.

- [Default RADIUS Configuration, page 1-27](#)
- [Identifying the RADIUS Server Host, page 1-27](#) (required)
- [Configuring RADIUS Login Authentication, page 1-30](#) (required)
- [Defining AAA Server Groups, page 1-32](#) (optional)
- [Configuring RADIUS Authorization for User Privileged Access and Network Services, page 1-34](#) (optional)
- [Starting RADIUS Accounting, page 1-35](#) (optional)
- [Configuring Settings for All RADIUS Servers, page 1-36](#) (optional)
- [Configuring the Switch to Use Vendor-Specific RADIUS Attributes, page 1-36](#) (optional)
- [Configuring the Switch for Vendor-Proprietary RADIUS Server Communication, page 1-38](#) (optional)
- [Configuring CoA on the Switch, page 1-39](#)
- [Monitoring and Troubleshooting CoA Functionality, page 1-40](#)
- [Configuring RADIUS Server Load Balancing, page 1-40](#) (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:

- Hostname or IP address
- Authentication destination port
- Accounting destination port
- Key string
- Timeout period
- Retransmission value

You identify RADIUS security servers by their hostname or IP address, hostname 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 `%RADIUS-4-RADIUS_DEAD` message appears, and then 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 settings on all RADIUS servers, see the [“Configuring Settings for All RADIUS Servers” section on page 1-36](#).

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 1-32](#).

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>]	<p>Specify the IP address or hostname of the remote RADIUS server host.</p> <ul style="list-style-type: none"> • (Optional) For auth-port <i>port-number</i>, specify the UDP destination port for authentication requests. • (Optional) For acct-port <i>port-number</i>, specify the UDP destination port for accounting requests. • (Optional) For timeout <i>seconds</i>, specify the time interval that the switch waits for the RADIUS server to reply before resending. The range is 1 to 1000. This setting overrides the radius-server timeout global configuration command setting. If no timeout is set with the radius-server host command, the setting of the radius-server timeout command is used. • (Optional) For retransmit <i>retries</i>, specify the number of times a RADIUS request is resent to a server if that server is not responding or responding slowly. The range is 1 to 1000. If no retransmit value is set with the radius-server host command, the setting of the radius-server retransmit global configuration command is used. • (Optional) For key <i>string</i>, specify the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. <p>Note The key is a text string that must match the encryption key used on the RADIUS server. Always configure the key as the last item in the radius-server host command. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.</p> <p>To configure the switch to recognize more than one host entry associated with a single IP address, enter this command as many times as necessary, making sure that each UDP port number is different. The switch software searches for hosts in the order in which you specify them. Set the timeout, retransmit, and encryption key values to use with the specific RADIUS host.</p>
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, see 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> ...]	<p>Create a login authentication method list.</p> <ul style="list-style-type: none"> To create a default list that is used when a named list is <i>not</i> specified in the login authentication command, use the default keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all 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. <p>Select one of these methods:</p> <ul style="list-style-type: none"> enable—Use the enable password for authentication. Before you can use this authentication method, you must define an enable password by using the enable password global configuration command. group radius—Use RADIUS authentication. Before you can use this authentication method, you must configure the RADIUS server. For more information, see the “Identifying the RADIUS Server Host” section on page 1-27. 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.
Step 5	login authentication { default <i>list-name</i> }	<p>Apply the authentication list to a line or set of lines.</p> <ul style="list-style-type: none"> If you specify default, use the default list created with the aaa authentication login command. For <i>list-name</i>, specify the list created with the aaa authentication login command.
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.

**Note**

To secure the switch for HTTP access by using AAA methods, you must configure the switch with the **ip http authentication aaa** global configuration command. Configuring AAA authentication does not secure the switch for HTTP access by using AAA methods.

For more information about the **ip http authentication** command, see the *Cisco IOS Security Command Reference, Release 12.4* on Cisco.com.

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>]	<p>Specify the IP address or hostname of the remote RADIUS server host.</p> <ul style="list-style-type: none"> • (Optional) For auth-port <i>port-number</i>, specify the UDP destination port for authentication requests. • (Optional) For acct-port <i>port-number</i>, specify the UDP destination port for accounting requests. • (Optional) For timeout <i>seconds</i>, specify the time interval that the switch waits for the RADIUS server to reply before resending. The range is 1 to 1000. This setting overrides the radius-server timeout global configuration command setting. If no timeout is set with the radius-server host command, the setting of the radius-server timeout command is used. • (Optional) For retransmit <i>retries</i>, specify the number of times a RADIUS request is resent to a server if that server is not responding or responding slowly. The range is 1 to 1000. If no retransmit value is set with the radius-server host command, the setting of the radius-server retransmit global configuration command is used. • (Optional) For key <i>string</i>, specify the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. <p>Note The key is a text string that must match the encryption key used on the RADIUS server. Always configure the key as the last item in the radius-server host command. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.</p> <p>To configure the switch to recognize more than one host entry associated with a single IP address, enter this command as many times as necessary, making sure that each UDP port number is different. The switch software searches for hosts in the order in which you specify them. Set the timeout, retransmit, and encryption key values to use with the specific RADIUS host.</p>
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.

	Command	Purpose
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 1-30.

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.

	Command	Purpose
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.

Establishing a Session with a Router if the AAA Server is Unreachable



Note

To configure this command, the switch must be running the LAN Base image.

The **aaa accounting system guarantee-first** command guarantees system accounting as the first record, which is the default condition. In some situations, users might be prevented from starting a session on the console or terminal connection until after the system reloads, which can take more than 3 minutes.

To establish a console or Telnet session with the router if the AAA server is unreachable when the router reloads, use the **no aaa accounting system guarantee-first** 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.
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 an interface 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 an interface 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, see 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]	<p>Enable the switch to recognize and use VSAs as defined by RADIUS IETF attribute 26.</p> <ul style="list-style-type: none"> (Optional) Use the accounting keyword to limit the set of recognized vendor-specific attributes to only accounting attributes. (Optional) Use the authentication keyword to limit the set of recognized vendor-specific attributes to only authentication attributes. <p>If you enter this command without keywords, both accounting and authentication vendor-specific attributes are used.</p>
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.



Note

For a complete list of RADIUS attributes or more information about vendor-specific attribute 26, see the "RADIUS Attributes" appendix in the *Cisco IOS Security Configuration Guide, Release 12.4*, on Cisco.com.

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 { <i>hostname</i> <i>ip-address</i> } non-standard	Specify the IP address or hostname of the remote RADIUS server host and identify that it is using a vendor-proprietary implementation of RADIUS.
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
```


Configuring CoA on the Switch

Beginning in privileged EXEC mode, follow these steps to configure CoA on a switch. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa new-model	Enable AAA.
Step 3	aaa server radius dynamic-author	Configure the switch as an authentication, authorization, and accounting (AAA) server to facilitate interaction with an external policy server.
Step 4	client { <i>ip-address</i> <i>name</i> } [vrf <i>vrfname</i>] [server-key <i>string</i>]	Enter dynamic authorization local server configuration mode and specify a RADIUS client from which a device will accept CoA and disconnect requests.
Step 5	server-key [0 7] <i>string</i>	Configure the RADIUS key to be shared between a device and RADIUS clients.
Step 6	port <i>port-number</i>	Specify the port on which a device listens for RADIUS requests from configured RADIUS clients.
Step 7	auth-type { any all session-key }	Specify the type of authorization the switch uses for RADIUS clients. The client must match all the configured attributes for authorization.
Step 8	ignore session-key	(Optional) Configure the switch to ignore the session-key. For more information about the ignore command, see the Cisco IOS Intelligent Services Gateway Command Reference on Cisco.com.
Step 9	ignore server-key	(Optional) Configure the switch to ignore the server-key. For more information about the ignore command, see the Cisco IOS Intelligent Services Gateway Command Reference on Cisco.com.
Step 10	authentication command bounce-port ignore	(Optional) Configure the switch to ignore a CoA request to temporarily disable the port hosting a session. The purpose of temporarily disabling the port is to trigger a DHCP renegotiation from the host when a VLAN change occurs and there is no supplicant on the endpoint to detect the change.
Step 11	authentication command disable-port ignore	(Optional) Configure the switch to ignore a nonstandard command requesting that the port hosting a session be administratively shut down. Shutting down the port results in termination of the session. Use standard CLI or SNMP commands to re-enable the port.
Step 12	end	Return to privileged EXEC mode.
Step 13	show running-config	Verify your entries.
Step 14	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 the AAA server functionality on the switch, use the **no aaa server radius dynamic authorization** global configuration command.

Monitoring and Troubleshooting CoA Functionality

Use these Cisco IOS commands to monitor and troubleshoot CoA functionality on the switch:

- **debug radius**
- **debug aaa coa**
- **debug aaa pod**
- **debug aaa subsys**
- **debug cmdhd [detail | error | events]**
- **show aaa attributes protocol radius**

Configuring RADIUS Server Load Balancing

This feature allows access and authentication requests to be evenly across all RADIUS servers in a server group. For more information, see the “RADIUS Server Load Balancing” chapter of the *Cisco IOS Security Configuration Guide*:

http://www.ciscosystems.com/en/US/docs/ios/12_2sb/feature/guide/sbrldbl.html

Displaying the RADIUS Configuration

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

Configuring the Switch for Local Authentication and Authorization

You can configure AAA to operate without a server by setting the switch to implement AAA in local mode. The switch then handles authentication and authorization. No accounting is available in this configuration.

Beginning in privileged EXEC mode, follow these steps to configure the switch for local AAA:

	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.

	Command	Purpose
Step 6	username <i>name</i> [privilege <i>level</i>] { password <i>encryption-type</i> <i>password</i> }	Enter the local database, and establish a username-based authentication system. Repeat this command for each user. <ul style="list-style-type: none"> For <i>name</i>, specify the user ID as one word. Spaces and quotation marks are not allowed. (Optional) For <i>level</i>, specify the privilege level the user has after gaining access. The range is 0 to 15. Level 15 gives privileged EXEC mode access. Level 0 gives user EXEC mode access. For <i>encryption-type</i>, enter 0 to specify that an unencrypted password follows. Enter 7 to specify that a hidden password follows. For <i>password</i>, specify the password the user must enter to gain access to the switch. The password must be from 1 to 25 characters, can contain embedded spaces, and must be the last option specified in the username command.
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.



Note

To secure the switch for HTTP access by using AAA methods, you must configure the switch with the **ip http authentication aaa** global configuration command. Configuring AAA authentication does not secure the switch for HTTP access by using AAA methods.

For more information about the **ip http authentication** command, see the *Cisco IOS Security Command Reference, Release 12.4*.

Configuring the Switch for Secure Shell

This section describes how to configure the Secure Shell (SSH) feature. To use this feature, you must install the cryptographic (encrypted) software image on your switch. You must obtain authorization to use this feature and to download the cryptographic software files from Cisco.com. For more information, see the release notes for this release.

- [Understanding SSH, page 1-42](#)
- [Configuring SSH, page 1-43](#)
- [Displaying the SSH Configuration and Status, page 1-45](#)

For SSH configuration examples, see the “SSH Configuration Examples” section in the “Configuring Secure Shell” chapter of the *Cisco IOS Security Configuration Guide*:

http://www.cisco.com/en/US/docs/ios/12_2/security/configuration/guide/scfssh.html

SSH functions the same in IPv6 as in IPv4. For IPv6, SSH supports IPv6 addresses and enables secure, encrypted connections with remote IPv6 nodes over an IPv6 transport.

**Note**

For complete syntax and usage information for the commands used in this section, see the command reference for this release and the command reference for Cisco IOS Release 12.2: http://www.cisco.com/en/US/docs/ios/12_2/security/command/reference/fsecur_r.html and the *Cisco IOS IPv6 Command Reference*.

Understanding SSH

SSH is a protocol that provides a secure, remote connection to a device. SSH provides more security for remote connections than Telnet does by providing strong encryption when a device is authenticated. This software release supports SSH Version 1 (SSHv1) and SSH Version 2 (SSHv2).

- [SSH Servers, Integrated Clients, and Supported Versions, page 1-42](#)
- [Limitations, page 1-42](#)

SSH Servers, Integrated Clients, and Supported Versions

The SSH feature has an SSH server and an SSH integrated client, which are applications that run on the switch. You can use an SSH client to connect to a switch running the SSH server. The SSH server works with the SSH client supported in this release and with non-Cisco SSH clients. The SSH client also works with the SSH server supported in this release and with non-Cisco SSH servers.

The switch supports an SSHv1 or an SSHv2 server.

The switch supports an SSHv1 client.

SSH supports the Data Encryption Standard (DES) encryption algorithm, the Triple DES (3DES) encryption algorithm, and password-based user authentication.

SSH also supports these user authentication methods:

- TACACS+ (for more information, see the [“Controlling Switch Access with TACACS+”](#) section on page 1-10)
- RADIUS (for more information, see the [“Controlling Switch Access with RADIUS”](#) section on page 1-17)
- Local authentication and authorization (for more information, see the [“Configuring the Switch for Local Authentication and Authorization”](#) section on page 1-40)

**Note**

The switch does not support IP Security (IPSec).

Limitations

These limitations apply to SSH:

- The switch supports Rivest, Shamir, and Adelman (RSA) authentication.
- SSH supports only the execution-shell application.
- The SSH server and the SSH client are supported only on DES (56-bit) and 3DES (168-bit) data encryption software.

- The switch supports the Advanced Encryption Standard (AES) encryption algorithm with a 128-bit key, 192-bit key, or 256-bit key. However, symmetric cipher AES to encrypt the keys is not supported.

Configuring SSH

This section has this configuration information:

- [Configuration Guidelines, page 1-43](#)
- [Setting Up the Switch to Run SSH, page 1-43](#) (required)
- [Configuring the SSH Server, page 1-44](#) (required only if you are configuring the switch as an SSH server)

Configuration Guidelines

Follow these guidelines when configuring the switch as an SSH server or SSH client:

- An RSA key pair generated by a SSHv1 server can be used by an SSHv2 server, and the reverse.
- If the SSH server is running on a stack master and the stack master fails, the new stack master uses the RSA key pair generated by the previous stack master.
- If you get CLI error messages after entering the **crypto key generate rsa** global configuration command, an RSA key pair has not been generated. Reconfigure the hostname and domain, and then enter the **crypto key generate rsa** command. For more information, see the “[Setting Up the Switch to Run SSH](#)” section on page 1-43.
- When generating the RSA key pair, the message `No host name specified` might appear. If it does, you must configure a hostname by using the **hostname** global configuration command.
- When generating the RSA key pair, the message `No domain specified` might appear. If it does, you must configure an IP domain name by using the **ip domain-name** global configuration command.
- When configuring the local authentication and authorization authentication method, make sure that AAA is disabled on the console.

Setting Up the Switch to Run SSH

Follow these steps to set up your switch to run SSH:

1. Download the cryptographic software image from Cisco.com. This step is required. For more information, see the release notes for this release.
2. Configure a hostname and IP domain name for the switch. Follow this procedure only if you are configuring the switch as an SSH server.
3. Generate an RSA key pair for the switch, which automatically enables SSH. Follow this procedure only if you are configuring the switch as an SSH server.
4. Configure user authentication for local or remote access. This step is required. For more information, see the “[Configuring the Switch for Local Authentication and Authorization](#)” section on page 1-40.

Beginning in privileged EXEC mode, follow these steps to configure a hostname and an IP domain name and to generate an RSA key pair. This procedure is required if you are configuring the switch as an SSH server.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	hostname <i>hostname</i>	Configure a hostname for your switch.
Step 3	ip domain-name <i>domain_name</i>	Configure a host domain for your switch.
Step 4	crypto key generate rsa	Enable the SSH server for local and remote authentication on the switch and generate an RSA key pair. We recommend that a minimum modulus size of 1024 bits. When you generate RSA keys, you are prompted to enter a modulus length. A longer modulus length might be more secure, but it takes longer to generate and to use.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip ssh or show ssh	Show the version and configuration information for your SSH server. Show the status of the SSH server on the switch.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete the RSA key pair, use the **crypto key zeroize rsa** global configuration command. After the RSA key pair is deleted, the SSH server is automatically disabled.

Configuring the SSH Server

Beginning in privileged EXEC mode, follow these steps to configure the SSH server:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip ssh version [1 2]	(Optional) Configure the switch to run SSH Version 1 or SSH Version 2. <ul style="list-style-type: none"> • 1—Configure the switch to run SSH Version 1. • 2—Configure the switch to run SSH Version 2. <p>If you do not enter this command or do not specify a keyword, the SSH server selects the latest SSH version supported by the SSH client. For example, if the SSH client supports SSHv1 and SSHv2, the SSH server selects SSHv2.</p>

	Command	Purpose
Step 3	ip ssh { <i>timeout seconds</i> authentication-retries <i>number</i> }	Configure the SSH control parameters: <ul style="list-style-type: none"> Specify the time-out value in seconds; the default is 120 seconds. The range is 0 to 120 seconds. This parameter applies to the SSH negotiation phase. After the connection is established, the switch uses the default time-out values of the CLI-based sessions. By default, up to five simultaneous, encrypted SSH connections for multiple CLI-based sessions over the network are available (session 0 to session 4). After the execution shell starts, the CLI-based session time-out value returns to the default of 10 minutes. <ul style="list-style-type: none"> Specify the number of times that a client can re-authenticate to the server. The default is 3; the range is 0 to 5. Repeat this step when configuring both parameters.
Step 4	line vty <i>line_number</i> [<i>ending_line_number</i>] transport input ssh	(Optional) Configure the virtual terminal line settings. <ul style="list-style-type: none"> Enter line configuration mode to configure the virtual terminal line settings. For <i>line_number</i> and <i>ending_line_number</i>, specify a pair of lines. The range is 0 to 15. Specify that the switch prevent non-SSH Telnet connections. This limits the router to only SSH connections.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip ssh or show ssh	Show the version and configuration information for your SSH server. Show the status of the SSH server connections on the switch.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default SSH control parameters, use the **no ip ssh {timeout | authentication-retries}** global configuration command.

Displaying the SSH Configuration and Status

To display the SSH server configuration and status, use one or more of the privileged EXEC commands in [Table 1-5](#):

Table 1-5 Commands for Displaying the SSH Server Configuration and Status

Command	Purpose
show ip ssh	Shows the version and configuration information for the SSH server.
show ssh	Shows the status of the SSH server.

For more information about these commands, see the “Secure Shell Commands” section in the “Other Security Features” chapter of the *Cisco IOS Security Command Reference*:

http://www.cisco.com/en/US/docs/ios/12_2/security/configuration/guide/scfpass.html

Configuring the Switch for Secure Socket Layer HTTP

This section describes how to configure Secure Socket Layer (SSL) version 3.0 support for the HTTP 1.1 server and client. SSL provides server authentication, encryption, and message integrity, as well as HTTP client authentication, to allow secure HTTP communications. To use this feature, the cryptographic (encrypted) software image must be installed on your switch. You must obtain authorization to use this feature and to download the cryptographic software files from Cisco.com. For more information about the crypto image, see the release notes for this release.

These sections contain this information:

- [Understanding Secure HTTP Servers and Clients, page 1-46](#)
- [Configuring Secure HTTP Servers and Clients, page 1-48](#)
- [Displaying Secure HTTP Server and Client Status, page 1-52](#)

For configuration examples and complete syntax and usage information for the commands used in this section, see the “HTTPS - HTTP Server and Client with SSL 3.0” feature description for Cisco IOS Release 12.2(15)T:

http://www.cisco.com/en/US/docs/ios/netmgmt/configuration/guide/nm_https_sc_ssl3.html

Understanding Secure HTTP Servers and Clients

On a secure HTTP connection, data to and from an HTTP server is encrypted before being sent over the Internet. HTTP with SSL encryption provides a secure connection to allow such functions as configuring a switch from a Web browser. Cisco's implementation of the secure HTTP server and secure HTTP client uses an implementation of SSL Version 3.0 with application-layer encryption. HTTP over SSL is abbreviated as HTTPS; the URL of a secure connection begins with `https://` instead of `http://`.

The primary role of the HTTP secure server (the switch) is to listen for HTTPS requests on a designated port (the default HTTPS port is 443) and pass the request to the HTTP 1.1 Web server. The HTTP 1.1 server processes requests and passes responses (pages) back to the HTTP secure server, which, in turn, responds to the original request.

The primary role of the HTTP secure client (the web browser) is to respond to Cisco IOS application requests for HTTPS User Agent services, perform HTTPS User Agent services for the application, and pass the response back to the application.

Certificate Authority Trustpoints

Certificate authorities (CAs) manage certificate requests and issue certificates to participating network devices. These services provide centralized security key and certificate management for the participating devices. Specific CA servers are referred to as *trustpoints*.

When a connection attempt is made, the HTTPS server provides a secure connection by issuing a certified X.509v3 certificate, obtained from a specified CA trustpoint, to the client. The client (usually a Web browser), in turn, has a public key that allows it to authenticate the certificate.

For secure HTTP connections, we highly recommend that you configure a CA trustpoint. If a CA trustpoint is not configured for the device running the HTTPS server, the server certifies itself and generates the needed RSA key pair. Because a self-certified (self-signed) certificate does not provide adequate security, the connecting client generates a notification that the certificate is self-certified, and the user has the opportunity to accept or reject the connection. This option is useful for internal network topologies (such as testing).

If you do not configure a CA trustpoint, when you enable a secure HTTP connection, either a temporary or a persistent self-signed certificate for the secure HTTP server (or client) is automatically generated.

- If the switch is not configured with a hostname and a domain name, a temporary self-signed certificate is generated. If the switch reboots, any temporary self-signed certificate is lost, and a new temporary new self-signed certificate is assigned.
- If the switch has been configured with a host and domain name, a persistent self-signed certificate is generated. This certificate remains active if you reboot the switch or if you disable the secure HTTP server so that it will be there the next time you re-enable a secure HTTP connection.

**Note**

The certificate authorities and trustpoints must be configured on each device individually. Copying them from other devices makes them invalid on the switch.

If a self-signed certificate has been generated, this information is included in the output of the **show running-config** privileged EXEC command. This is a partial sample output from that command displaying a self-signed certificate.

```
Switch# show running-config
Building configuration...

<output truncated>

crypto pki trustpoint TP-self-signed-3080755072
  enrollment selfsigned
  subject-name cn=IOS-Self-Signed-Certificate-3080755072
  revocation-check none
  rsakeypair TP-self-signed-3080755072
!
!
crypto ca certificate chain TP-self-signed-3080755072
certificate self-signed 01
  3082029F 30820208 A0030201 02020101 300D0609 2A864886 F70D0101 04050030
  59312F30 2D060355 04031326 494F532D 53656C66 2D536967 6E65642D 43657274
  69666963 6174652D 33303830 37353530 37323126 30240609 2A864886 F70D0109
```

<output truncated>

You can remove this self-signed certificate by disabling the secure HTTP server and entering the **no crypto pki trustpoint TP-self-signed-30890755072** global configuration command. If you later re-enable a secure HTTP server, a new self-signed certificate is generated.

**Note**

The values that follow *TP self-signed* depend on the serial number of the device.

You can use an optional command (**ip http secure-client-auth**) to allow the HTTPS server to request an X.509v3 certificate from the client. Authenticating the client provides more security than server authentication by itself.

For additional information on Certificate Authorities, see the “Configuring Certification Authority Interoperability” chapter in the *Cisco IOS Security Configuration Guide, Release 124* on Cisco.com.

CipherSuites

A CipherSuite specifies the encryption algorithm and the digest algorithm to use on a SSL connection. When connecting to the HTTPS server, the client Web browser offers a list of supported CipherSuites, and the client and server negotiate the best encryption algorithm to use from those on the list that are supported by both. For example, Netscape Communicator 4.76 supports U.S. security with RSA Public Key Cryptography, MD2, MD5, RC2-CBC, RC4, DES-CBC, and DES-EDE3-CBC.

For the best possible encryption, you should use a client browser that supports 128-bit encryption, such as Microsoft Internet Explorer Version 5.5 (or later) or Netscape Communicator Version 4.76 (or later). The SSL_RSA_WITH_DES_CBC_SHA CipherSuite provides less security than the other CipherSuites, as it does not offer 128-bit encryption.

The more secure and more complex CipherSuites require slightly more processing time. This list defines the CipherSuites supported by the switch and ranks them from fastest to slowest in terms of router processing load (speed):

1. SSL_RSA_WITH_DES_CBC_SHA—RSA key exchange (RSA Public Key Cryptography) with DES-CBC for message encryption and SHA for message digest
2. SSL_RSA_WITH_RC4_128_MD5—RSA key exchange with RC4 128-bit encryption and MD5 for message digest
3. SSL_RSA_WITH_RC4_128_SHA—RSA key exchange with RC4 128-bit encryption and SHA for message digest
4. SSL_RSA_WITH_3DES_EDE_CBC_SHA—RSA key exchange with 3DES and DES-EDE3-CBC for message encryption and SHA for message digest

RSA (in conjunction with the specified encryption and digest algorithm combinations) is used for both key generation and authentication on SSL connections. This usage is independent of whether or not a CA trustpoint is configured.

Configuring Secure HTTP Servers and Clients

- [Default SSL Configuration, page 1-48](#)
- [SSL Configuration Guidelines, page 1-48](#)
- [Configuring a CA Trustpoint, page 1-49](#)
- [Configuring the Secure HTTP Server, page 1-50](#)
- [Configuring the Secure HTTP Client, page 1-51](#)

Default SSL Configuration

The standard HTTP server is enabled.

SSL is enabled.

No CA trustpoints are configured.

No self-signed certificates are generated.

SSL Configuration Guidelines

When SSL is used in a switch cluster, the SSL session terminates at the cluster commander. Cluster member switches must run standard HTTP.

Before you configure a CA trustpoint, you should ensure that the system clock is set. If the clock is not set, the certificate is rejected due to an incorrect date.

In a switch stack, the SSL session terminates at the stack master.

Configuring a CA Trustpoint

For secure HTTP connections, we recommend that you configure an official CA trustpoint. A CA trustpoint is more secure than a self-signed certificate.

Beginning in privileged EXEC mode, follow these steps to configure a CA trustpoint:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	hostname <i>hostname</i>	Specify the hostname of the switch (required only if you have not previously configured a hostname). The hostname is required for security keys and certificates.
Step 3	ip domain-name <i>domain-name</i>	Specify the IP domain name of the switch (required only if you have not previously configured an IP domain name). The domain name is required for security keys and certificates.
Step 4	crypto key generate rsa	(Optional) Generate an RSA key pair. RSA key pairs are required before you can obtain a certificate for the switch. RSA key pairs are generated automatically. You can use this command to regenerate the keys, if needed.
Step 5	crypto ca trustpoint <i>name</i>	Specify a local configuration name for the CA trustpoint and enter CA trustpoint configuration mode.
Step 6	enrollment url <i>url</i>	Specify the URL to which the switch should send certificate requests.
Step 7	enrollment http-proxy <i>host-name</i> <i>port-number</i>	(Optional) Configure the switch to obtain certificates from the CA through an HTTP proxy server.
Step 8	crl query <i>url</i>	Configure the switch to request a certificate revocation list (CRL) to ensure that the certificate of the peer has not been revoked.
Step 9	primary	(Optional) Specify that the trustpoint should be used as the primary (default) trustpoint for CA requests.
Step 10	exit	Exit CA trustpoint configuration mode and return to global configuration mode.
Step 11	crypto ca authentication <i>name</i>	Authenticate the CA by getting the public key of the CA. Use the same name used in Step 5.
Step 12	crypto ca enroll <i>name</i>	Obtain the certificate from the specified CA trustpoint. This command requests a signed certificate for each RSA key pair.
Step 13	end	Return to privileged EXEC mode.
Step 14	show crypto ca trustpoints	Verify the configuration.
Step 15	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no crypto ca trustpoint** *name* global configuration command to delete all identity information and certificates associated with the CA.

Configuring the Secure HTTP Server

If you are using a certificate authority for certification, you should use the previous procedure to configure the CA trustpoint on the switch before enabling the HTTP server. If you have not configured a CA trustpoint, a self-signed certificate is generated the first time that you enable the secure HTTP server. After you have configured the server, you can configure options (path, access list to apply, maximum number of connections, or timeout policy) that apply to both standard and secure HTTP servers.

Beginning in privileged EXEC mode, follow these steps to configure a secure HTTP server:

	Command	Purpose
Step 1	show ip http server status	(Optional) Display the status of the HTTP server to determine if the secure HTTP server feature is supported in the software. You should see one of these lines in the output: HTTP secure server capability: Present or HTTP secure server capability: Not present
Step 2	configure terminal	Enter global configuration mode.
Step 3	ip http secure-server	Enable the HTTPS server if it has been disabled. The HTTPS server is enabled by default.
Step 4	ip http secure-port <i>port-number</i>	(Optional) Specify the port number to be used for the HTTPS server. The default port number is 443. Valid options are 443 or any number in the range 1025 to 65535.
Step 5	ip http secure-ciphersuite { [3des-ede-cbc-sha] [rc4-128-md5] [rc4-128-sha] [des-cbc-sha]}	(Optional) Specify the CipherSuites (encryption algorithms) to be used for encryption over the HTTPS connection. If you do not have a reason to specify a particularly CipherSuite, you should allow the server and client to negotiate a CipherSuite that they both support. This is the default.
Step 6	ip http secure-client-auth	(Optional) Configure the HTTP server to request an X.509v3 certificate from the client for authentication during the connection process. The default is for the client to request a certificate from the server, but the server does not attempt to authenticate the client.
Step 7	ip http secure-trustpoint <i>name</i>	Specify the CA trustpoint to use to get an X.509v3 security certificate and to authenticate the client certificate connection. Note Use of this command assumes you have already configured a CA trustpoint according to the previous procedure.
Step 8	ip http path <i>path-name</i>	(Optional) Set a base HTTP path for HTML files. The path specifies the location of the HTTP server files on the local system (usually located in system flash memory).
Step 9	ip http access-class <i>access-list-number</i>	(Optional) Specify an access list to use to allow access to the HTTP server.
Step 10	ip http max-connections <i>value</i>	(Optional) Set the maximum number of concurrent connections that are allowed to the HTTP server. The range is 1 to 16; the default value is 5.

	Command	Purpose
Step 11	ip http timeout-policy <i>idle seconds life seconds requests value</i>	(Optional) Specify how long a connection to the HTTP server can remain open under the defined circumstances: <ul style="list-style-type: none"> idle—the maximum time period when no data is received or response data cannot be sent. The range is 1 to 600 seconds. The default is 180 seconds (3 minutes). life—the maximum time period from the time that the connection is established. The range is 1 to 86400 seconds (24 hours). The default is 180 seconds. requests—the maximum number of requests processed on a persistent connection. The maximum value is 86400. The default is 1.
Step 12	end	Return to privileged EXEC mode.
Step 13	show ip http server secure status	Display the status of the HTTP secure server to verify the configuration.
Step 14	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip http server** global configuration command to disable the standard HTTP server. Use the **no ip http secure-server** global configuration command to disable the secure HTTP server. Use the **no ip http secure-port** and the **no ip http secure-ciphersuite** global configuration commands to return to the default settings. Use the **no ip http secure-client-auth** global configuration command to remove the requirement for client authentication.

To verify the secure HTTP connection by using a Web browser, enter `https://URL`, where the URL is the IP address or hostname of the server switch. If you configure a port other than the default port, you must also specify the port number after the URL. For example:

```
https://209.165.129:1026
or
https://host.domain.com:1026
```

Configuring the Secure HTTP Client

The standard HTTP client and secure HTTP client are always enabled. A certificate authority is required for secure HTTP client certification. This procedure assumes that you have previously configured a CA trustpoint on the switch. If a CA trustpoint is not configured and the remote HTTPS server requires client authentication, connections to the secure HTTP client fail.

Beginning in privileged EXEC mode, follow these steps to configure a secure HTTP client:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip http client secure-trustpoint <i>name</i>	(Optional) Specify the CA trustpoint to be used if the remote HTTP server requests client authentication. Using this command assumes that you have already configured a CA trustpoint by using the previous procedure. The command is optional if client authentication is not needed or if a primary trustpoint has been configured.
Step 3	ip http client secure-ciphersuite {[3des-ede-cbc-sha] [rc4-128-md5] [rc4-128-sha] [des-cbc-sha]}	(Optional) Specify the CipherSuites (encryption algorithms) to be used for encryption over the HTTPS connection. If you do not have a reason to specify a particular CipherSuite, you should allow the server and client to negotiate a CipherSuite that they both support. This is the default.

	Command	Purpose
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip http client secure status	Display the status of the HTTP secure server to verify the configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip http client secure-trustpoint** *name* to remove a client trustpoint configuration. Use the **no ip http client secure-ciphersuite** to remove a previously configured CipherSuite specification for the client.

Displaying Secure HTTP Server and Client Status

To display the SSL secure server and client status, use the privileged EXEC commands in [Table 1-6](#):

Table 1-6 Commands for Displaying the SSL Secure Server and Client Status

Command	Purpose
show ip http client secure status	Shows the HTTP secure client configuration.
show ip http server secure status	Shows the HTTP secure server configuration.
show running-config	Shows the generated self-signed certificate for secure HTTP connections.

Configuring the Switch for Secure Copy Protocol

The Secure Copy Protocol (SCP) feature provides a secure and authenticated method for copying switch configurations or switch image files. SCP relies on Secure Shell (SSH), an application and a protocol that provides a secure replacement for the Berkeley r-tools.

For SSH to work, the switch needs an RSA public/private key pair. This is the same with SCP, which relies on SSH for its secure transport.

Because SSH also relies on AAA authentication, and SCP relies further on AAA authorization, correct configuration is necessary.

- Before enabling SCP, you must correctly configure SSH, authentication, and authorization on the switch.
- Because SCP relies on SSH for its secure transport, the router must have an Rivest, Shamir, and Adelman (RSA) key pair.



Note

When using SCP, you cannot enter the password into the copy command. You must enter the password when prompted.

Information About Secure Copy

To configure the Secure Copy feature, you should understand these concepts.

The behavior of SCP is similar to that of remote copy (rcp), which comes from the Berkeley r-tools suite, except that SCP relies on SSH for security. SCP also requires that authentication, authorization, and accounting (AAA) authorization be configured so the router can determine whether the user has the correct privilege level.

A user who has appropriate authorization can use SCP to copy any file in the Cisco IOS File System (IFS) to and from a switch by using the **copy** command. An authorized administrator can also do this from a workstation.

For information about how to configure and verify SCP, see the “Secure Copy Protocol” section in the *Cisco IOS Security Configuration Guide: Securing User Services, Release 12.4*:
http://www.cisco.com/en/US/docs/ios/sec_user_services/configuration/guide/sec_secure_copy_ps6350_TSD_Products_Configuration_Guide_Chapter.html



CHAPTER 1

Configuring IEEE 802.1x Port-Based Authentication

IEEE 802.1x port-based authentication prevents unauthorized devices (clients) from gaining access to the network. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

The Catalyst 2960 switch command reference and the “RADIUS Commands” section in the Cisco IOS Security Command Reference, Release 12.4, have command syntax and usage information.

This chapter includes these sections:

- [Understanding IEEE 802.1x Port-Based Authentication, page 1-1](#)
- [Configuring 802.1x Authentication, page 1-34](#)
- [Displaying 802.1x Statistics and Status, page 1-66](#)

Understanding IEEE 802.1x Port-Based Authentication

The standard defines a client-server-based access control and authentication protocol to prevent 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 switch or LAN services.

Until the client is authenticated, IEEE 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, normal traffic passes through the port.

- [Device Roles, page 1-3](#)
- [Authentication Process, page 1-4](#)
- [Authentication Initiation and Message Exchange, page 1-5](#)
- [Authentication Manager, page 1-7](#)
- [Ports in Authorized and Unauthorized States, page 1-10](#)
- [802.1x Authentication and Switch Stacks, page 1-11](#)
- [802.1x Host Mode, page 1-12](#)

- [Multidomain Authentication, page 1-12](#)
- [802.1x Multiple Authentication Mode, page 1-13](#)
- [MAC Move, page 1-14](#)
- [MAC Replace, page 1-15](#)
- [802.1x Accounting, page 1-16](#)
- [802.1x Accounting Attribute-Value Pairs, page 1-16](#)
- [802.1x Readiness Check, page 1-17](#)
- [802.1x Authentication with VLAN Assignment, page 1-17](#)
- [Using 802.1x Authentication with Per-User ACLs, page 1-19](#)
- [802.1x Authentication with Guest VLAN, page 1-22](#)
- [802.1x Authentication with Restricted VLAN, page 1-23](#)
- [802.1x Authentication with Inaccessible Authentication Bypass, page 1-24](#)



Note To use IEEE 802.1x authentication with inaccessible authentication bypass, the switch must be running the LAN base image.

- [802.1x Authentication with Voice VLAN Ports, page 1-26](#)
- [802.1x Authentication with Port Security, page 1-27](#)
- [802.1x Authentication with Wake-on-LAN, page 1-27](#)
- [802.1x Authentication with MAC Authentication Bypass, page 1-27](#)
- [802.1x User Distribution, page 1-29](#)
- [Network Admission Control Layer 2 802.1x Validation, page 1-29](#)



Note To use Network Admission Control, the switch must be running the LAN base image.


- [Flexible Authentication Ordering, page 1-30](#)
- [Open1x Authentication, page 1-30](#)
- [Using Voice Aware 802.1x Security, page 1-31](#)
- [802.1x Supplicant and Authenticator Switches with Network Edge Access Topology \(NEAT\), page 1-31](#)
- [802.1x Authentication with Downloadable ACLs and Redirect URLs, page 1-20](#)
- [Using IEEE 802.1x Authentication with ACLs and the RADIUS Filter-Id Attribute, page 1-32](#)



Note To use IEEE 802.1x authentication with ACLs and the Filter-Id attribute, the switch must be running the LAN base image.

- [Common Session ID, page 1-33](#)

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 802.1x standard.)



Note To resolve Windows XP network connectivity and 802.1x authentication issues, read the Microsoft Knowledge Base article:
<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. (The switch is the *authenticator* in the 802.1x standard.)

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 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-E, Catalyst 3560-E, Catalyst 3750, Catalyst 3560, Catalyst 3550, Catalyst 2975, Catalyst 2970, Catalyst 2960, 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.

Authentication Process

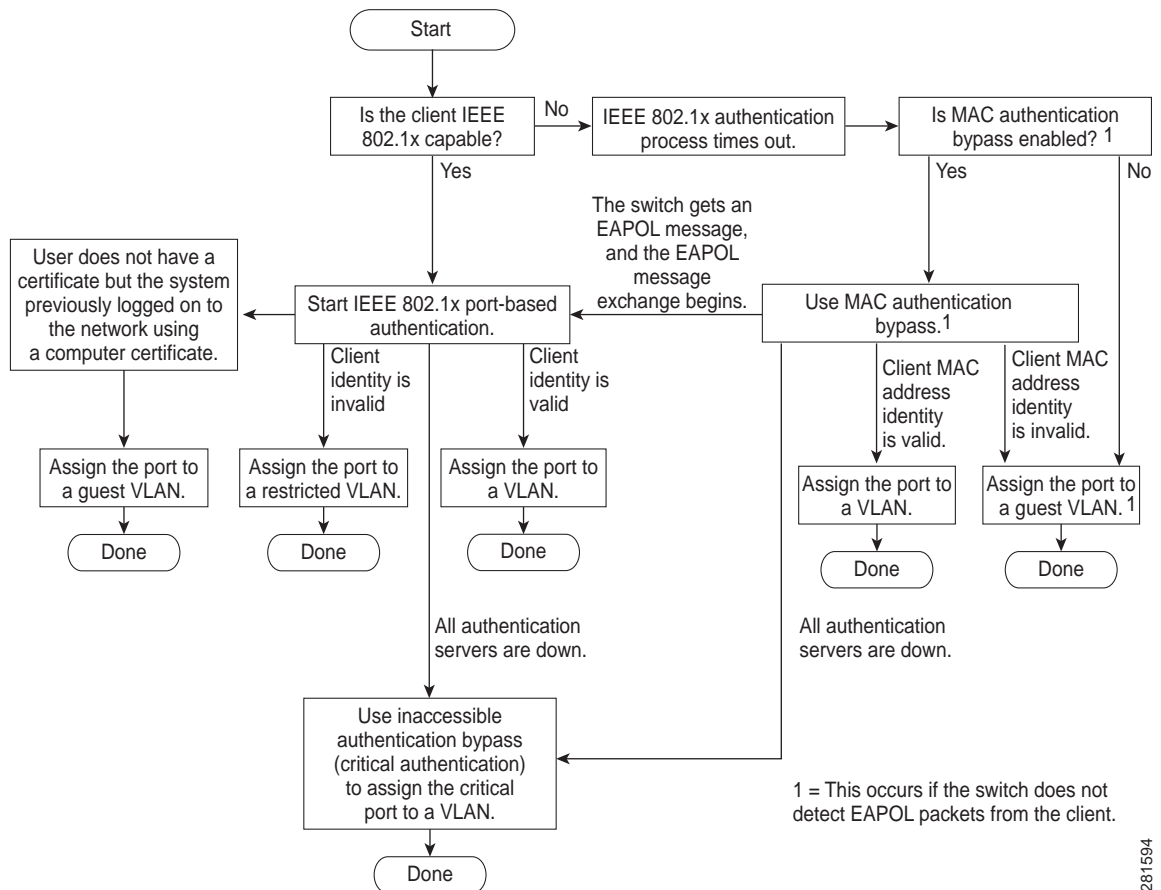
When 802.1x port-based authentication is enabled and the client supports 802.1x-compliant client software, these events occur:

- If the client identity is valid and the 802.1x authentication succeeds, the switch grants the client access to the network.
- If 802.1x authentication times out while waiting for an EAPOL message exchange and MAC authentication bypass is enabled, the switch can use the client MAC address for authorization. If the client MAC address is valid and the authorization succeeds, the switch grants the client access to the network. If the client MAC address is invalid and the authorization fails, the switch assigns the client to a guest VLAN that provides limited services if a guest VLAN is configured.
- If the switch gets an invalid identity from an 802.1x-capable client and a restricted VLAN is specified, the switch can assign the client to a restricted VLAN that provides limited services.
- If the RADIUS authentication server is unavailable (down) and inaccessible authentication bypass is enabled, the switch grants the client access to the network by putting the port in the critical-authentication state in the RADIUS-configured or the user-specified access VLAN.



Note Inaccessible authentication bypass is also referred to as critical authentication or the AAA fail policy.

Figure 1-2 Authentication Flowchart



281594

The switch re-authenticates a client when one of these situations occurs:

- Periodic re-authentication is enabled, and the re-authentication timer expires.

You can configure the re-authentication timer to use a switch-specific value or to be based on values from the RADIUS server.

After 802.1x authentication using a RADIUS server is configured, the switch uses timers based on the Session-Timeout RADIUS attribute (Attribute[27]) and the Termination-Action RADIUS attribute (Attribute [29]).

The Session-Timeout RADIUS attribute (Attribute[27]) specifies the time after which re-authentication occurs.

The Termination-Action RADIUS attribute (Attribute [29]) specifies the action to take during re-authentication. The actions are *Initialize* and *ReAuthenticate*. When the *Initialize* action is set (the attribute value is *DEFAULT*), the 802.1x session ends, and connectivity is lost during re-authentication. When the *ReAuthenticate* action is set (the attribute value is RADIUS-Request), the session is not affected during re-authentication.

- You manually re-authenticate the client by entering the **dot1x re-authenticate interface interface-id** privileged EXEC command.

If Multidomain authentication (MDA) is enabled on a port, this flow can be used with some exceptions that are applicable to voice authorization. For more information on MDA, see the “[Multidomain Authentication](#)” section on page 1-12.

Authentication Initiation and Message Exchange

During 802.1x authentication, the switch or the client can initiate authentication. If you enable authentication on a port by using the **authentication port-control auto** interface configuration command, the switch initiates authentication when the link state changes from down to up or periodically as long as the port remains up and unauthenticated. The switch sends an EAP-request/identity frame to the client to request its identity. Upon receipt of the frame, the client responds with an EAP-response/identity frame.

However, if during boot up, 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.



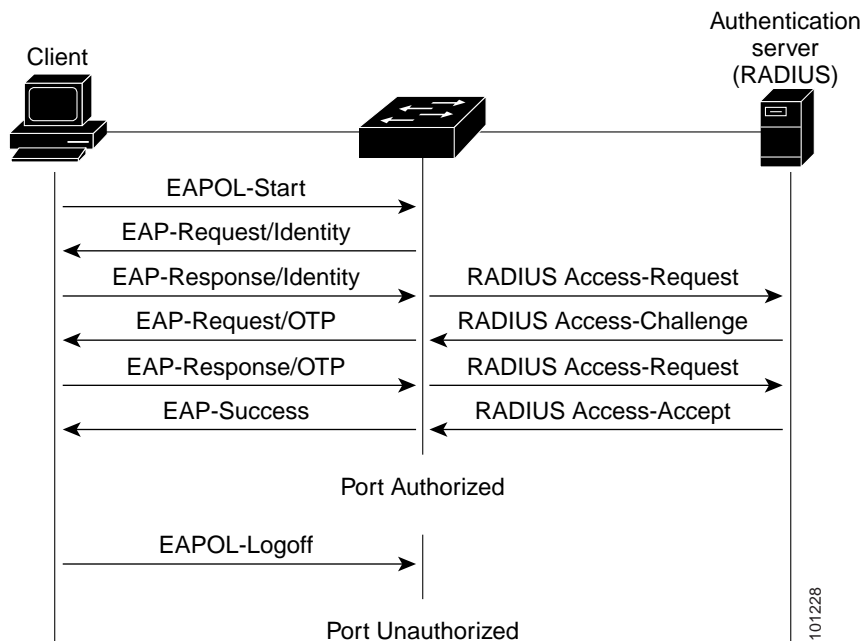
Note

If 802.1x authentication 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 1-10.

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. If the authentication fails, authentication can be retried, the port might be assigned to a VLAN that provides limited services, or network access is not granted. For more information, see the “[Ports in Authorized and Unauthorized States](#)” section on page 1-10.

The specific exchange of EAP frames depends on the authentication method being used. [Figure 1-3](#) shows a message exchange initiated by the client when the client uses the One-Time-Password (OTP) authentication method with a RADIUS server.

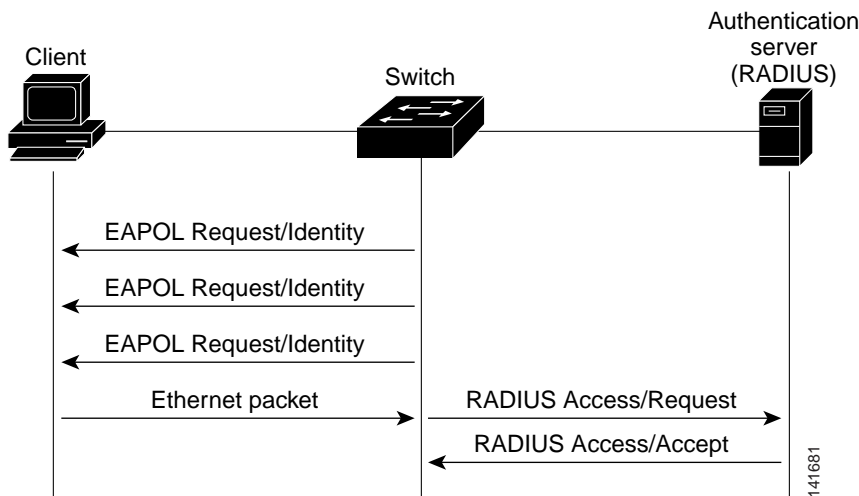
Figure 1-3 Message Exchange



If 802.1x authentication times out while waiting for an EAPOL message exchange and MAC authentication bypass is enabled, the switch can authorize the client when the switch detects an Ethernet packet from the client. The switch uses the MAC address of the client as its identity and includes this information in the RADIUS-access/request frame that is sent to the RADIUS server. After the server sends the switch the RADIUS-access/accept frame (authorization is successful), the port becomes authorized. If authorization fails and a guest VLAN is specified, the switch assigns the port to the guest VLAN. If the switch detects an EAPOL packet while waiting for an Ethernet packet, the switch stops the MAC authentication bypass process and stops 802.1x authentication.

Figure 1-4 shows the message exchange during MAC authentication bypass.

Figure 1-4 Message Exchange During MAC Authentication Bypass



Authentication Manager

In Cisco IOS Release 12.2(46)SE and earlier, you could not use the same authorization methods, including CLI commands and messages, on this switch and also on other network devices, such as a Catalyst 6000. You had to use separate authentication configurations. Cisco IOS Release 12.2(50)SE and later supports the same authorization methods on all Catalyst switches in a network.

Cisco IOS Release 12.2(55)SE supports filtering verbose system messages from the authentication manager. For details, see the “[Authentication Manager CLI Commands](#)” section on page 1-9.

- [Port-Based Authentication Methods, page 1-7](#)
- [Per-User ACLs and Filter-Ids, page 1-8](#)
- [Authentication Manager CLI Commands, page 1-9](#)

Port-Based Authentication Methods

Table 1-1 lists the authentication methods supported in these host modes:

- Single host—Only one data or voice host (client) can be authenticated on a port.
- Multiple host—Multiple data hosts can be authenticated on the same port. (If a port becomes unauthorized in multiple-host mode, the switch denies network access to all of the attached clients.)
- Multidomain authentication (MDA) —Both a data device and voice device can be authenticated on the same switch port. The port is divided into a data domain and a voice domain.
- Multiple authentication—Multiple hosts can authenticate on the data VLAN. This mode also allows one client on the VLAN if a voice VLAN is configured.

Table 1-1 802.1x Features

Authentication method	Mode			
	Single Host	Multiple Host	MDA ¹	Multiple Authentication ²
802.1x	VLAN assignment Per-user ACL Filter-ID attribute Downloadable ACL ³ Redirect URL ³	VLAN assignment Per-user ACL Filter-ID attribute Downloadable ACL ⁴ Redirect URL ³	VLAN assignment Per-user ACL ³ Filter-ID attribute ³ Downloadable ACL ³ Redirect URL ³	Per-user ACL ³ Filter-Id attribute ³ Downloadable ACL ³ Redirect URL ³
MAC authentication bypass	VLAN assignment Per-user ACL Filter-ID attribute Downloadable ACL ³ Redirect URL ³	VLAN assignment Per-user ACL Filter-ID attribute Downloadable ACL ³ Redirect URL ³	VLAN assignment Per-user ACL ³ Filter-ID attribute ³ Downloadable ACL ³ Redirect URL ³	Per-user ACL ³ Filter-Id attribute ³ Downloadable ACL ³ Redirect URL ³
Standalone web authentication ⁴	Proxy ACL, Filter-Id attribute, downloadable ACL ²			

Table 1-1 802.1x Features (continued)

Authentication method	Mode			
	Single Host	Multiple Host	MDA ¹	Multiple Authentication ²
NAC Layer 2 IP validation	Filter-Id attribute ³	Filter-Id attribute ³	Filter-Id attribute ³	Filter-Id attribute ³
	Downloadable ACL	Downloadable ACL	Downloadable ACL	Downloadable ACL ³
	Redirect URL	Redirect URL	Redirect URL	Redirect URL ³
Web authentication as fallback method ⁵	Proxy ACL	Proxy ACL	Proxy ACL	Proxy ACL ³
	Filter-Id attribute ³	Filter-Id attribute ³	Filter-Id attribute ³	Filter-Id attribute ³
	Downloadable ACL ³	Downloadable ACL ³	Downloadable ACL ³	Downloadable ACL ³

1. MDA = Multidomain authentication.
2. Also referred to as *multiauth*.
3. Supported in Cisco IOS Release 12.2(50)SE and later.
4. Supported in Cisco IOS Release 12.2(50)SE and later.
5. For clients that do not support 802.1x authentication.

Per-User ACLs and Filter-Ids

In releases earlier than Cisco IOS Release 12.2(50)SE, per-user ACLs and filter Ids were only supported in single-host mode. In Cisco IOS Release 12.2(50), support was added for MDA- and multiauth-enabled ports. In 12.2(52)SE and later, support was added for ports in multihost mode.

In releases earlier than Cisco IOS Release 12.2(50)SE, an ACL configured on the switch is not compatible with an ACL configured on another device running Cisco IOS software, such as a Catalyst 6000 switch.

In Cisco IOS Release 12.2(50)SE or later, the ACLs configured on the switch are compatible with other devices running the Cisco IOS release.



Note

You can only set **any** as the source in the ACL.



Note

For any ACL configured for multiple-host mode, the source portion of statement must be *any*. (For example, **permit icmp any host 10.10.1.1**.)

You must specify *any* in the source ports of any defined ACL. Otherwise, the ACL cannot be applied and authorization fails. Single host is the only exception to support backward compatibility.

More than one host can be authenticated on MDA-enabled and multiauth ports. The ACL policy applied for one host does not effect the traffic of another host.

If only one host is authenticated on a multi-host port, and the other hosts gain network access without authentication, the ACL policy for the first host can be applied to the other connected hosts by specifying *any* in the source address.

Authentication Manager CLI Commands

The authentication-manager interface-configuration commands control all the authentication methods, such as 802.1x, MAC authentication bypass, and web authentication. The authentication manager commands determine the priority and order of authentication methods applied to a connected host.

The authentication manager commands control generic authentication features, such as host-mode, violation mode, and the authentication timer. Generic authentication commands include the **authentication host-mode**, **authentication violation**, and **authentication timer** interface configuration commands.

802.1x-specific commands begin with the **dot1x** or **authentication** keyword. For example, the **authentication port-control auto** interface configuration command enables authentication on an interface. However, the **dot1x system-auth-control** global configuration command only globally enables or disables 802.1x authentication.



Note

If 802.1x authentication is globally disabled, other authentication methods are still enabled on that port, such as web authentication.

The **authentication manager** commands provide the same functionality as earlier 802.1x commands.

Table 1-2 Authentication Manager Commands and Earlier 802.1x Commands

The authentication manager commands in Cisco IOS Release 12.2(50)SE or later	The equivalent 802.1x commands in Cisco IOS Release 12.2(46)SE and earlier	Description
authentication control-direction {both in}	dot1x control-direction {both in}	Enable authentication with the wake-on-LAN (WoL) feature, and configure the port control as unidirectional or bidirectional.
authentication event	dot1x auth-fail vlan dot1x critical (interface configuration) dot1x guest-vlan6	Enable the restricted VLAN on a port. Enable the inaccessible-authentication-bypass feature. Specify an active VLAN as a guest VLAN.
authentication fallback <i>fallback-profile</i>	dot1x fallback <i>fallback-profile</i>	Configure a port to use web authentication as a fallback method for clients that do not support authentication.
authentication host-mode [multi-auth multi-domain multi-host single-host]	dot1x host-mode {single-host multi-host multi-domain}	Allow a single host (client) or multiple hosts on an authorized port.
authentication order	dot1x mac-auth-bypass	Provides the flexibility to define the order of authentication methods to be used.
authentication periodic	dot1x reauthentication	Enable periodic re-authentication of the client.
authentication port-control {auto force-authorized force-unauthorized}	dot1x port-control {auto force-authorized force-unauthorized}	Enable manual control of the authorization state of the port.

Table 1-2 Authentication Manager Commands and Earlier 802.1x Commands (continued)

The authentication manager commands in Cisco IOS Release 12.2(50)SE or later	The equivalent 802.1x commands in Cisco IOS Release 12.2(46)SE and earlier	Description
<code>authentication timer</code>	<code>dot1x timeout</code>	Set the timers.
<code>authentication violation {protect restrict shutdown}</code>	<code>dot1x violation-mode {shutdown restrict protect}</code>	Configure the violation modes that occur when a new device connects to a port or when a new device connects to a port after the maximum number of devices are connected to that port.

Beginning with Cisco IOS Release 12.2(55)SE, you can filter out verbose system messages generated by the authentication manager. The filtered content typically relates to authentication success. You can also filter verbose messages for 802.1x authentication and MAB authentication. There is a separate command for each authentication method:

- The **no authentication logging verbose** global configuration command filters verbose messages from the authentication manager.
- The **no dot1x logging verbose** global configuration command filters 802.1x authentication verbose messages.
- The **no mab logging verbose** global configuration command filters MAC authentication bypass (MAB) verbose messages

For more information, see the command reference for this release.

Ports in Authorized and Unauthorized States

During 802.1x authentication, 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 that is not configured as a voice VLAN port disallows all ingress and egress traffic except for 802.1x authentication, CDP, and STP packets. When a client is successfully authenticated, the port changes to the *authorized* state, allowing all traffic for the client to flow normally. If the port is configured as a voice VLAN port, the port allows VoIP traffic and 802.1x protocol packets before the client is successfully authenticated.

If a client that does not support 802.1x authentication connects 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 standard, 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 **authentication port-control** interface configuration command and these keywords:

- **force-authorized**—disables 802.1x authentication and causes the port to change 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 changes 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 change to the unauthorized state.

If the link state of a port changes from up to down, or if an EAPOL-logoff frame is received, the port returns to the unauthorized state.

802.1x Authentication and Switch Stacks



Note

Switch stacks are supported only on Catalyst 2960-S switches running the LAN base image.

If a switch is added to or removed from a switch stack, 802.1x authentication is not affected as long as the IP connectivity between the RADIUS server and the stack remains intact. This statement also applies if the stack master is removed from the switch stack. Note that if the stack master fails, a stack member becomes the new stack master by using the election process described in [Chapter 1, “Managing Switch Stacks,”](#) and the 802.1x authentication process continues as usual.

If IP connectivity to the RADIUS server is interrupted because the switch that was connected to the server is removed or fails, these events occur:

- Ports that are already authenticated and that do not have periodic re-authentication enabled remain in the authenticated state. Communication with the RADIUS server is not required.
- Ports that are already authenticated and that have periodic re-authentication enabled (with the **dot1x re-authentication** global configuration command) fail the authentication process when the re-authentication occurs. Ports return to the unauthenticated state during the re-authentication process. Communication with the RADIUS server is required.

For an ongoing authentication, the authentication fails immediately because there is no server connectivity.

If the switch that failed comes up and rejoins the switch stack, the authentications might or might not fail depending on the boot-up time and whether the connectivity to the RADIUS server is re-established by the time the authentication is attempted.


To avoid loss of connectivity to the RADIUS server, you should ensure that there is a redundant connection to it. For example, you can have a redundant connection to the stack master and another to a stack member, and if the stack master fails, the switch stack still has connectivity to the RADIUS server.

802.1x Host Mode

You can configure an 802.1x port for single-host or for multiple-hosts mode. In single-host mode (see [Figure 1-1 on page 1-3](#)), only one client can be connected to the 802.1x-enabled switch port. The switch detects the client by sending an EAPOL frame 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.

In multiple-hosts mode, you can attach multiple hosts to a single 802.1x-enabled port. [Figure 1-5 on page 1-12](#) shows 802.1x port-based authentication in a wireless LAN. In this mode, only one of the attached clients must be authorized for all clients to be granted network access. If the port becomes unauthorized (re-authentication fails or an EAPOL-logoff message is received), the switch denies network access to all of the attached clients. In this topology, the wireless access point is responsible for authenticating the clients attached to it, and it also acts as a client to the switch.

Figure 1-5 Multiple Host Mode Example



The switch supports multidomain authentication (MDA), which allows both a data device and a voice device, such as an IP Phone (Cisco or non-Cisco), to connect to the same switch port. For more information, see the [“Multidomain Authentication” section on page 1-12](#).

Multidomain Authentication

The switch supports multidomain authentication (MDA), which allows both a data device and voice device, such as an IP phone (Cisco or non-Cisco), to authenticate on the same switch port. The port is divided into a data domain and a voice domain.



Note To use MDA, the switch must be running the LAN base image.

MDA does not enforce the order of device authentication. However, for best results, we recommend that a voice device is authenticated before a data device on an MDA-enabled port.

Follow these guidelines for configuring MDA:

- To configure a switch port for MDA, see the [“Configuring the Host Mode” section on page 1-44](#).
- You must configure the voice VLAN for the IP phone when the host mode is set to multidomain. For more information, see [Chapter 1, “Configuring VLANs.”](#)
- Voice VLAN assignment on an MDA-enabled port is supported in Cisco IOS Release 12.2(40)SE and later.



Note If you use a dynamic VLAN to assign a voice VLAN on an MDA-enabled switch port on a switch running Cisco IOS Release 12.2(37)SE, the voice device fails authorization.

- To authorize a voice device, the AAA server must be configured to send a Cisco Attribute-Value (AV) pair attribute with a value of `device-traffic-class=voice`. Without this value, the switch treats the voice device as a data device.
- The guest VLAN and restricted VLAN features only apply to the data devices on an MDA-enabled port. The switch treats a voice device that fails authorization as a data device.
- If more than one device attempts authorization on either the voice or the data domain of a port, it is error disabled.
- Until a device is authorized, the port drops its traffic. Non-Cisco IP phones or voice devices are allowed into both the data and voice VLANs. The data VLAN allows the voice device to contact a DHCP server to obtain an IP address and acquire the voice VLAN information. After the voice device starts sending on the voice VLAN, its access to the data VLAN is blocked.
- A voice device MAC address that is binding on the data VLAN is not counted towards the port security MAC address limit.
- MDA can use MAC authentication bypass as a fallback mechanism to allow the switch port to connect to devices that do not support 802.1x authentication. For more information, see the [“MAC Authentication Bypass” section on page 1-37](#).
- When a *data* or a *voice* device is detected on a port, its MAC address is blocked until authorization succeeds. If the authorization fails, the MAC address remains blocked for 5 minutes.
- If more than five devices are detected on the *data* VLAN or more than one voice device is detected on the *voice* VLAN while a port is unauthorized, the port is error disabled.
- When a port host mode changes from single- or multihost to multidomain mode, an authorized data device remains authorized on the port. However, a Cisco IP phone on the port voice VLAN is automatically removed and must be reauthenticated on that port.
- Active fallback mechanisms such as guest VLAN and restricted VLAN remain configured after a port changes from single-host or multiple-host mode to multidomain mode.
- Switching a port host mode from multidomain to single-host or multiple-hosts mode removes all authorized devices from the port.
- If a data domain is authorized first and placed in the guest VLAN, non-802.1x-capable voice devices need their packets tagged on the voice VLAN to trigger authentication. The phone need not need to send tagged traffic. (The same is true for an 802.1x-capable phone.)
- We do not recommend per-user ACLs with an MDA-enabled port. An authorized device with a per-user ACL policy might impact traffic on both the port voice and data VLANs. You can use only one device on the port to enforce per-user ACLs.

For more information, see the [“Configuring the Host Mode” section on page 1-44](#).

802.1x Multiple Authentication Mode

Multiple-authentication (multiauth) mode allows multiple authenticated clients on the data VLAN. Each host is individually authenticated. If a voice VLAN is configured, this mode also allows one client on the VLAN. (If the port detects any additional voice clients, they are discarded from the port, but no violation errors occur.)

If a hub or access point is connected to an 802.1x-enabled port, each connected client must be authenticated.

For non-802.1x devices, you can use MAC authentication bypass or web authentication as the per-host authentication fallback method to authenticate different hosts with different methods on a single port.

There is no limit to the number of data hosts that can authenticate on a multi-auth port. However, only one voice device is allowed if the voice VLAN is configured. Since there is no host limit defined, a violation will not be triggered, if a second voice device is seen, we silently discard it but do not trigger a violation.

For MDA functionality on the voice VLAN, multiple-authentication mode assigns authenticated devices to either a data or a voice VLAN, depending on the VSAs received from the authentication server.

**Note**

When a port is in multiple-authentication mode, the guest VLAN and the authentication-failed VLAN features do not activate.

For more information about critical authentication mode and the critical VLAN, see the [“802.1x Authentication with Inaccessible Authentication Bypass” section on page 1-24](#).

For more information about configuring multi-auth mode on a port, see the [“Configuring the Host Mode” section on page 1-44](#).

Beginning with Cisco IOS Release 12.2(55)SE, you can assign a RADIUS-server-supplied VLAN in multi-auth mode, under these conditions:

- The switch is running the LAN base image.
- The host is the first host authorized on the port, and the RADIUS server supplies VLAN information.
- Subsequent hosts are authorized with a VLAN that matches the operational VLAN.
- A host is authorized on the port with no VLAN assignment, and subsequent hosts either have no VLAN assignment, or their VLAN information matches the operational VLAN.
- The first host authorized on the port has a group VLAN assignment, and subsequent hosts either have no VLAN assignment, or their group VLAN matches the group VLAN on the port. Subsequent hosts must use the same VLAN from the VLAN group as the first host. If a VLAN list is used, all hosts are subject to the conditions specified in the VLAN list.
- Only one voice VLAN assignment is supported on a multi-auth port.
- After a VLAN is assigned to a host on the port, subsequent hosts must have matching VLAN information or be denied access to the port.
- You cannot configure a guest VLAN or an auth-fail VLAN in multi-auth mode.
- The behavior of the critical-auth VLAN is not changed for multi-auth mode. When a host tries to authenticate and the server is not reachable, all authorized hosts are reinitialized in the configured VLAN.

MAC Move

When a MAC address is authenticated on one switch port, that address is not allowed on another authentication manager-enabled port of the switch. If the switch detects that same MAC address on another authentication manager-enabled port, the address is not allowed.

There are situations where a MAC address might need to move from one port to another on the same switch. For example, when there is another device (for example a hub or an IP phone) between an authenticated host and a switch port, you might want to disconnect the host from the device and connect it directly to another port on the same switch.

You can globally enable MAC move so the device is reauthenticated on the new port. When a host moves to a second port, the session on the first port is deleted, and the host is reauthenticated on the new port.

MAC move is supported on all host modes. (The authenticated host can move to any port on the switch, no matter which host mode is enabled on the that port.)

When a MAC address moves from one port to another, the switch terminates the authenticated session on the original port and initiates a new authentication sequence on the new port.

The MAC move feature applies to both voice and data hosts.

**Note**

In open authentication mode, a MAC address is immediately moved from the original port to the new port, with no requirement for authorization on the new port.

For more information see the [“Enabling MAC Move” section on page 1-49](#).

MAC Replace

**Note**

To configure the MAC replace feature, the switch must be running the LAN base image.

Beginning with Cisco IOS Release 12.2(55)SE, the MAC replace feature can be configured to address the violation that occurs when a host attempts to connect to a port where another host was previously authenticated.

**Note**

This feature does not apply to ports in multi-auth mode, because violations are not triggered in that mode. It does not apply to ports in multiple host mode, because in that mode, only the first host requires authentication.

If you configure the **authentication violation** interface configuration command with the **replace** keyword, the authentication process on a port in multi-domain mode is:

- A new MAC address is received on a port with an existing authenticated MAC address.
- The authentication manager replaces the MAC address of the current data host on the port with the new MAC address.
- The authentication manager initiates the authentication process for the new MAC address.
- If the authentication manager determines that the new host is a voice host, the original voice host is removed.

If a port is in open authentication mode, any new MAC address is immediately added to the MAC address table.

For more information see the [“Enabling MAC Replace” section on page 1-49](#).

802.1x Accounting

The 802.1x standard defines how users are authorized and authenticated for network access but does not keep track of network usage. 802.1x accounting is disabled by default. You can enable 802.1x accounting to monitor this activity on 802.1x-enabled ports:

- User successfully authenticates.
- User logs off.
- Link-down occurs.
- Re-authentication successfully occurs.
- Re-authentication fails.

The switch does not log 802.1x accounting information. Instead, it sends this information to the RADIUS server, which must be configured to log accounting messages.

802.1x Accounting Attribute-Value Pairs

The information sent to the RADIUS server is represented in the form of Attribute-Value (AV) pairs. These AV pairs provide data for different applications. (For example, a billing application might require information that is in the Acct-Input-Octets or the Acct-Output-Octets attributes of a RADIUS packet.)

AV pairs are automatically sent by a switch that is configured for 802.1x accounting. Three types of RADIUS accounting packets are sent by a switch:

- START—sent when a new user session starts
- INTERIM—sent during an existing session for updates
- STOP—sent when a session terminates

Table 1-3 lists the AV pairs and when they are sent by the switch:

Table 1-3 Accounting AV Pairs

Attribute Number	AV Pair Name	START	INTERIM	STOP
Attribute[1]	User-Name	Always	Always	Always
Attribute[4]	NAS-IP-Address	Always	Always	Always
Attribute[5]	NAS-Port	Always	Always	Always
Attribute[8]	Framed-IP-Address	Never	Sometimes ¹	Sometimes ¹
Attribute[25]	Class	Always	Always	Always
Attribute[30]	Called-Station-ID	Always	Always	Always
Attribute[31]	Calling-Station-ID	Always	Always	Always
Attribute[40]	Acct-Status-Type	Always	Always	Always
Attribute[41]	Acct-Delay-Time	Always	Always	Always
Attribute[42]	Acct-Input-Octets	Never	Always	Always
Attribute[43]	Acct-Output-Octets	Never	Always	Always
Attribute[44]	Acct-Session-ID	Always	Always	Always
Attribute[45]	Acct-Authentic	Always	Always	Always

Table 1-3 Accounting AV Pairs (continued)

Attribute Number	AV Pair Name	START	INTERIM	STOP
Attribute[46]	Acct-Session-Time	Never	Always	Always
Attribute[49]	Acct-Terminate-Cause	Never	Never	Always
Attribute[61]	NAS-Port-Type	Always	Always	Always

1. The Framed-IP-Address AV pair is sent only if a valid Dynamic Host Control Protocol (DHCP) binding exists for the host in the DHCP snooping bindings table.

You can view the AV pairs that are being sent by the switch by entering the **debug radius accounting** privileged EXEC command. For more information about this command, see the *Cisco IOS Debug Command Reference*:

http://www.cisco.com/en/US/docs/ios/12_2/debug/command/reference/122debug.html

For more information about AV pairs, see RFC 3580, “802.1x Remote Authentication Dial In User Service (RADIUS) Usage Guidelines.”

802.1x Readiness Check



Note

To use 802.1x readiness check, the switch must be running the LAN base image.

The 802.1x readiness check monitors 802.1x activity on all the switch ports and displays information about the devices connected to the ports that support 802.1x. You can use this feature to determine if the devices connected to the switch ports are 802.1x-capable. You use an alternate authentication such as MAC authentication bypass or web authentication for the devices that do not support 802.1x functionality.

This feature only works if the supplicant on the client supports a query with the NOTIFY EAP notification packet. The client must respond within the 802.1x timeout value.

For information on configuring the switch for the 802.1x readiness check, see the “[Configuring 802.1x Readiness Check](#)” section on page 1-38.

802.1x Authentication with VLAN Assignment

The RADIUS server sends the VLAN assignment to configure the switch port. The RADIUS server database maintains the username-to-VLAN mappings, assigning 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.

Voice device authentication is supported with multidomain host mode in Cisco IOS Release 12.2(37)SE. In Cisco IOS Release 12.2(40)SE and later, when a voice device is authorized and the RADIUS server returned an authorized VLAN, the voice VLAN on the port is configured to send and receive packets on the assigned voice VLAN. Voice VLAN assignment behaves the same as data VLAN assignment on multidomain authentication (MDA)-enabled ports. For more information, see the “[Multidomain Authentication](#)” section on page 1-12.

When configured on the switch and the RADIUS server, 802.1x authentication with VLAN assignment has these characteristics:

- If no VLAN is supplied by the RADIUS server or if 802.1x authentication is disabled, the port is configured in its access VLAN after successful authentication. Recall that an access VLAN is a VLAN assigned to an access port. All packets sent from or received on this port belong to this VLAN.
- If 802.1x authentication is enabled but the VLAN information from the RADIUS server is not valid, authorization fails and configured VLAN remains in use. This prevents ports from appearing unexpectedly in an inappropriate VLAN because of a configuration error.

Configuration errors could include specifying a malformed VLAN ID, a nonexistent VLAN ID, an RSPAN VLAN, a shut down or suspended VLAN. In the case of a multidomain host port, configuration errors can also be due to an attempted assignment of a data VLAN that matches the configured or assigned voice VLAN ID (or the reverse).

- If 802.1x authentication is enabled and all information from the RADIUS server is valid, the authorized device 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.
- Enabling port security does not impact the RADIUS server-assigned VLAN behavior.
- If 802.1x authentication is disabled on the port, it is returned to the configured access VLAN and configured voice 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. In the case of a multidomain host, the same applies to voice devices when the port is fully authorized with these exceptions:
 - If the VLAN configuration change of one device results in matching the other device configured or assigned VLAN, then authorization of all devices on the port is terminated and multidomain host mode is disabled until a valid configuration is restored where data and voice device configured VLANs no longer match.
 - If a voice device is authorized and is using a downloaded voice VLAN, the removal of the voice VLAN configuration, or modifying the configuration value to *dot1p* or *untagged* results in voice device un-authorization and the disablement of multi-domain host mode.

When the port is in the force authorized, force unauthorized, unauthorized, or shutdown state, it is put into the configured access VLAN.

The 802.1x authentication 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 interface configuration from the RADIUS server.
- Enable 802.1x authentication. (The VLAN assignment feature is automatically enabled when you configure 802.1x authentication 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, VLAN ID, or VLAN-Group
 - [83] Tunnel-Preference

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 1-36](#).

Using 802.1x Authentication 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 on the same switch. However, a port ACL takes precedence over a router ACL. If you apply input port ACL to an interface 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 `inacl#<n>` for the ingress direction and `outacl#<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 1, “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).

The maximum size of the per-user ACL is 4000 ASCII characters but is limited by the maximum size of RADIUS-server per-user ACLs.

For examples of vendor-specific attributes, see the [“Configuring the Switch to Use Vendor-Specific RADIUS Attributes” section on page 1-36](#). For more information about configuring ACLs, see [Chapter 1, “Configuring Network Security with ACLs.”](#)



Note

Per-user ACLs are supported only in single-host mode.

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

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

For more configuration information, see the [“Authentication Manager” section on page 1-7](#).

802.1x Authentication with Downloadable ACLs and Redirect URLs

You can download ACLs and redirect URLs from a RADIUS server to the switch during 802.1x authentication or MAC authentication bypass of the host. You can also download ACLs during web authentication.



Note A downloadable ACL is also referred to as a *dACL*.

If more than one host is authenticated and the host is in single-host, MDA, or multiple-authentication mode, the switch changes the source address of the ACL to the host IP address.

You can apply the ACLs and redirect URLs to all the devices connected to the 802.1x-enabled port.

If no ACLs are downloaded during 802.1x authentication, the switch applies the static default ACL on the port to the host. On a voice VLAN port configured in multi-auth or MDA mode, the switch applies the ACL only to the phone as part of the authorization policies.

Beginning with Cisco IOS Release 12.2(55)SE, if there is no static ACL on a port, a dynamic auth-default ACL is created, and policies are enforced before dACLs are downloaded and applied.



Note The auth-default-ACL is created only when the switch is running the LAN base image.



Note The auth-default-ACL does not appear in the running configuration.

The auth-default ACL is created when at least one host with an authorization policy is detected on the port. The auth-default ACL is removed from the port when the last authenticated session ends. You can configure the auth-default ACL by using the **ip access-list extended auth-default-acl** global configuration command.



Note The auth-default-ACL does not support Cisco Discovery Protocol (CDP) bypass in the single host mode. You must configure a static ACL on the interface to support CDP bypass.

The 802.1x and MAB authentication methods support two authentication modes, *open* and *closed*. If there is no static ACL on a port in *closed* authentication mode:

- An auth-default-ACL is created.
- The auth-default-ACL allows only DHCP traffic until policies are enforced.

- When the first host authenticates, the authorization policy is applied without IP address insertion.
- When a second host is detected, the policies for the first host are refreshed, and policies for the first and subsequent sessions are enforced with IP address insertion.

If there is no static ACL on a port in *open* authentication mode:

- An auth-default-ACL-OPEN is created and allows all traffic.
- Policies are enforced with IP address insertion to prevent security breaches.
- Web authentication is subject to the auth-default-ACL-OPEN.

To control access for hosts with no authorization policy, you can configure a directive. The supported values for the directive are *open* and *default*. When you configure the *open* directive, all traffic is allowed. The *default* directive subjects traffic to the access provided by the port. You can configure the directive either in the user profile on the AAA server or on the switch. To configure the directive on the AAA server, use the **authz-directive =<open/default>** global command. To configure the directive on the switch, use the **epm access-control open** global configuration command.



Note

The default value of the directive is *default*.

If a host falls back to web authentication on a port without a configured ACL:

- If the port is in open authentication mode, the auth-default-ACL-OPEN is created.
- If the port is in closed authentication mode, the auth-default-ACL is created.

The access control entries (ACEs) in the fallback ACL are converted to per-user entries. If the configured fallback profile does not include a fallback ACL, the host is subject to the auth-default-ACL associated with the port.



Note

If you use a custom logo with web authentication and it is stored on an external server, the port ACL must allow access to the external server before authentication. You must either configure a static port ACL or change the auth-default-ACL to provide appropriate access to the external server.

Cisco Secure ACS and Attribute-Value Pairs for the Redirect URL

The switch uses these *cisco-av-pair* VSAs:

- url-redirect is the HTTP to HTTPS URL.
- url-redirect-acl is the switch ACL name or number.

The switch uses the CiscoSecure-Defined-ACL attribute value pair to intercept an HTTP or HTTPS request from the end point device. The switch then forwards the client web browser to the specified redirect address. The url-redirect attribute value pair on the Cisco Secure ACS contains the URL to which the web browser is redirected. The url-redirect-acl attribute value pair contains the name or number of an ACL that specifies the HTTP or HTTPS traffic to redirect. Traffic that matches a permit ACE in the ACL is redirected.



Note

Define the URL redirect ACL and the default port ACL on the switch.

If a redirect URL is configured for a client on the authentication server, a default port ACL on the connected client switch port must also be configured.

Cisco Secure ACS and Attribute-Value Pairs for Downloadable ACLs

You can set the CiscoSecure-Defined-ACL Attribute-Value pair on the Cisco Secure ACS with the RADIUS cisco-av-pair vendor-specific attributes (VSAs). This pair specifies the names of the downloadable ACLs on the Cisco Secure ACS with the #ACL#-IP-name-number attribute.

- The *name* is the ACL name.
- The *number* is the version number (for example, 3f783768).

If a downloadable ACL is configured for a client on the authentication server, a default port ACL on the connected client switch port must also be configured.

If the default ACL is configured on the switch and the Cisco Secure ACS sends a host-access-policy to the switch, it applies the policy to traffic from the host connected to a switch port. If the policy does not apply, the switch applies the default ACL. If the Cisco Secure ACS sends the switch a downloadable ACL, this ACL takes precedence over the default ACL that is configured on the switch port. However, if the switch receives an host access policy from the Cisco Secure ACS but the default ACL is not configured, the authorization failure is declared.

For configuration details, see the [“Authentication Manager” section on page 1-7](#) and the [“Configuring 802.1x Authentication with Downloadable ACLs and Redirect URLs” section on page 1-60](#).

VLAN ID-based MAC Authentication

You can use VLAN ID-based MAC authentication if you wish to authenticate hosts based on a static VLAN ID instead of a downloadable VLAN. When you have a static VLAN policy configured on your switch, VLAN information is sent to an IAS (Microsoft) RADIUS server along with the MAC address of each host for authentication. The VLAN ID configured on the connected port is used for MAC authentication. By using VLAN ID-based MAC authentication with an IAS server, you can have a fixed number of VLANs in the network.

The feature also limits the number of VLANs monitored and handled by STP. The network can be managed as a fixed VLAN.



Note

This feature is not supported on Cisco ACS Server. (The ACS server ignores the sent VLAN-IDs for new hosts and only authenticates based on the MAC address.)

For configuration information, see the [“Configuring VLAN ID-based MAC Authentication” section on page 1-63](#). Additional configuration is similar MAC authentication bypass, as described in the [“Configuring MAC Authentication Bypass” section on page 1-56](#).

802.1x Authentication with Guest VLAN

You can configure a guest VLAN for each 802.1x port on the switch to provide limited services to clients, such as downloading 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 you enable a guest VLAN on an 802.1x port, the switch assigns clients to a guest VLAN when the switch does not receive a response to its EAP request/identity frame or when EAPOL packets are not sent by the client. The port is automatically set to multi-host mode.

The switch maintains the EAPOL packet history. If an EAPOL packet is detected on the interface during the lifetime of the link, the switch determines that the device connected to that interface is an 802.1x-capable supplicant, and the interface does not change to the guest VLAN state. EAPOL history is cleared if the interface link status goes down. If no EAPOL packet is detected on the interface, the interface changes to the guest VLAN state.

If devices send EAPOL packets to the switch during the lifetime of the link, the switch no longer allows clients that fail authentication access to the guest VLAN.

If the switch is trying to authorize an 802.1x-capable voice device and the AAA server is unavailable, the authorization attempt fails, but the detection of the EAPOL packet is saved in the EAPOL history. When the AAA server becomes available, the switch authorizes the voice device. However, the switch no longer allows other devices access to the guest VLAN. To prevent this situation, use one of these command sequences:

- Enter the **authentication event no-response action authorize vlan** *vlan-id* interface configuration command to allow access to the guest VLAN.
- Enter the **shutdown** interface configuration command followed by the **no shutdown** interface configuration command to restart the port.

**Note**

If an EAPOL packet is detected after the interface has changed to the guest VLAN, the interface reverts to an unauthorized state, and 802.1x authentication restarts.

Any number of 802.1x-incapable clients are allowed access when the switch port is moved to the guest VLAN. If an 802.1x-capable client 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, multiple host, or multi-domain modes.

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 trunk ports; it is supported only on access ports.

The switch supports *MAC authentication bypass*. When MAC authentication bypass is enabled on an 802.1x port, the switch can authorize clients based on the client MAC address when 802.1x authentication times out while waiting for an EAPOL message exchange. After detecting a client on an 802.1x port, the switch waits for an Ethernet packet from the client. The switch sends the authentication server a RADIUS-access/request frame with a username and password based on the MAC address. If authorization succeeds, the switch grants the client access to the network. If authorization fails, the switch assigns the port to the guest VLAN if one is specified. For more information, see the [“802.1x Authentication with MAC Authentication Bypass” section on page 1-27](#).

For more information, see the [“Configuring a Guest VLAN” section on page 1-51](#).

802.1x Authentication with Restricted VLAN

You can configure a restricted VLAN (also referred to as an *authentication failed VLAN*) for each 802.1x port on a switch stack or a switch to provide limited services to clients that cannot access the guest VLAN. These clients are 802.1x-compliant and cannot access another VLAN because they fail the authentication process. A restricted VLAN allows users without valid credentials in an authentication server (typically, visitors to an enterprise) to access a limited set of services. The administrator can control the services available to the restricted VLAN.

**Note**

You can configure a VLAN to be both the guest VLAN and the restricted VLAN if you want to provide the same services to both types of users.

Without this feature, the client attempts and fails authentication indefinitely, and the switch port remains in the spanning-tree blocking state. With this feature, you can configure the switch port to be in the restricted VLAN after a specified number of authentication attempts (the default value is 3 attempts).

The authenticator counts the failed authentication attempts for the client. When this count exceeds the configured maximum number of authentication attempts, the port moves to the restricted VLAN. The failed attempt count increments when the RADIUS server replies with either an *EAP failure* or an empty response without an EAP packet. When the port moves into the restricted VLAN, the failed attempt counter resets.

Users who fail authentication remain in the restricted VLAN until the next re-authentication attempt. A port in the restricted VLAN tries to re-authenticate at configured intervals (the default is 60 seconds). If re-authentication fails, the port remains in the restricted VLAN. If re-authentication is successful, the port moves either to the configured VLAN or to a VLAN sent by the RADIUS server. You can disable re-authentication. If you do this, the only way to restart the authentication process is for the port to receive a *link down* or *EAP logoff* event. We recommend that you keep re-authentication enabled if a client might connect through a hub. When a client disconnects from the hub, the port might not receive the *link down* or *EAP logoff* event.

After a port moves to the restricted VLAN, a simulated EAP success message is sent to the client. This prevents clients from indefinitely attempting authentication. Some clients (for example, devices running Windows XP) cannot implement DHCP without EAP success.

Restricted VLANs are supported only on 802.1x ports in single-host mode and on Layer 2 ports.

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

Other security features such as dynamic ARP Inspection, DHCP snooping, and IP source guard can be configured independently on a restricted VLAN.

For more information, see the [“Configuring a Restricted VLAN” section on page 1-52](#).

802.1x Authentication with Inaccessible Authentication Bypass

Use the inaccessible authentication bypass feature, also referred to as *critical authentication* or the *AAA fail policy*, when the switch cannot reach the configured RADIUS servers and new hosts cannot be authenticated. You can configure the switch to connect those hosts to *critical ports*.

When a new host tries to connect to the critical port, that host is moved to a user-specified access VLAN, the *critical VLAN*. The administrator gives limited authentication to the hosts.

When the switch tries to authenticate a host connected to a critical port, the switch checks the status of the configured RADIUS server. If a server is available, the switch can authenticate the host. However, if all the RADIUS servers are unavailable, the switch grants network access to the host and puts the port in the *critical-authentication* state, which is a special case of the authentication state.

Support on Multiple-Authentication Ports

When a port is configured on any host mode and the AAA server is unavailable, the port is then configured to multi-host mode and moved to the critical VLAN. To support this inaccessible bypass on multiple-authentication (multiauth) ports, use the **authentication event server dead action reinitialize vlan *vlan-id*** command. When a new host tries to connect to the critical port, that port is reinitialized and all the connected hosts are moved to the user-specified access VLAN.

This command is supported on all host modes.

Authentication Results

The behavior of the inaccessible authentication bypass feature depends on the authorization state of the port:

- If the port is unauthorized when a host connected to a critical port tries to authenticate and all servers are unavailable, the switch puts the port in the critical-authentication state in the RADIUS-configured or user-specified access VLAN.
- If the port is already authorized and reauthentication occurs, the switch puts the critical port in the critical-authentication state in the current VLAN, which might be the one previously assigned by the RADIUS server.
- If the RADIUS server becomes unavailable during an authentication exchange, the current exchange times out, and the switch puts the critical port in the critical-authentication state during the next authentication attempt.

You can configure the critical port to reinitialize hosts and move them out of the critical VLAN when the RADIUS server is again available. When this is configured, all critical ports in the critical-authentication state are automatically re-authenticated. For more information, see the command reference for this release and the [“Configuring the Inaccessible Authentication Bypass Feature”](#) section on page 1-54.

Feature Interactions

Inaccessible authentication bypass interacts with these features:

- Guest VLAN—Inaccessible authentication bypass is compatible with guest VLAN. When a guest VLAN is enabled on 802.1x port, the features interact as follows:
 - If at least one RADIUS server is available, the switch assigns a client to a guest VLAN when the switch does not receive a response to its EAP request/identity frame or when EAPOL packets are not sent by the client.
 - If all the RADIUS servers are not available and the client is connected to a critical port, the switch authenticates the client and puts the critical port in the critical-authentication state in the RADIUS-configured or user-specified access VLAN.
 - If all the RADIUS servers are not available and the client is not connected to a critical port, the switch might not assign clients to the guest VLAN if one is configured.
 - If all the RADIUS servers are not available and if a client is connected to a critical port and was previously assigned to a guest VLAN, the switch keeps the port in the guest VLAN.
- Restricted VLAN—If the port is already authorized in a restricted VLAN and the RADIUS servers are unavailable, the switch puts the critical port in the critical-authentication state in the restricted VLAN.
- 802.1x accounting—Accounting is not affected if the RADIUS servers are unavailable.

- Private VLAN—You can configure inaccessible authentication bypass on a private VLAN host port. The access VLAN must be a secondary private VLAN.
- Voice VLAN—Inaccessible authentication bypass is compatible with voice VLAN, but the RADIUS-configured or user-specified access VLAN and the voice VLAN must be different.
- Remote Switched Port Analyzer (RSPAN)—Do not configure an RSPAN VLAN as the RADIUS-configured or user-specified access VLAN for inaccessible authentication bypass.

In a switch stack:



Note

Switch stacks are supported only on Catalyst 2960-S switches running the LAN base image.

- The stack master checks the status of the RADIUS servers by sending keepalive packets. When the status of a RADIUS server changes, the stack master sends the information to the stack members. The stack members can then check the status of RADIUS servers when re-authenticating critical ports.
- If the new stack master is elected, the link between the switch stack and RADIUS server might change, and the new stack immediately sends keepalive packets to update the status of the RADIUS servers.

If the server status changes from *dead* to *alive*, the switch re-authenticates all switch ports in the critical-authentication state.
- When a member is added to the stack, the stack master sends the member the server status.

802.1x Authentication 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.

The IP phone uses the VVID for its voice traffic, regardless of the authorization state of the port. This allows the phone to work independently of 802.1x authentication.

In single-host mode, only the IP phone is allowed on the voice VLAN. In multiple-hosts mode, additional clients can send traffic on the voice VLAN after a supplicant is authenticated on the PVID. When multiple-hosts mode is enabled, the supplicant authentication affects both the PVID and the VVID.

A voice VLAN port becomes active when there is a 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 authentication is enabled on a voice VLAN port, the switch drops packets from unrecognized IP phones more than one hop away.

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

When IP phones are connected to an 802.1x-enabled switch port that is in single host mode, the switch grants the phones network access without authenticating them. We recommend that you use multidomain authentication (MDA) on the port to authenticate both a data device and a voice device, such as an IP phone.

**Note**

If you enable 802.1x authentication on an access port on which a voice VLAN is configured and to which a Cisco IP Phone is connected, the Cisco IP phone loses connectivity to the switch for up to 30 seconds.

For more information about voice VLANs, see [Chapter 1, “Configuring Voice VLAN.”](#)

802.1x Authentication with Port Security

In general, Cisco does not recommend enabling port security when IEEE 802.1x is enabled. Since IEEE 802.1x enforces a single MAC address per port (or per VLAN when MDA is configured for IP telephony), port security is redundant and in some cases may interfere with expected IEEE 802.1x operations.

802.1x Authentication with Wake-on-LAN

The 802.1x authentication with the wake-on-LAN (WoL) feature allows dormant PCs to be powered when the switch receives a specific Ethernet frame, known as the *magic packet*. You can use this feature in environments where administrators need to connect to systems that have been powered down.

When a host that uses WoL is attached through an 802.1x port and the host powers off, the 802.1x port becomes unauthorized. The port can only receive and send EAPOL packets, and WoL magic packets cannot reach the host. When the PC is powered off, it is not authorized, and the switch port is not opened.

When the switch uses 802.1x authentication with WoL, the switch forwards traffic to unauthorized 802.1x ports, including magic packets. While the port is unauthorized, the switch continues to block ingress traffic other than EAPOL packets. The host can receive packets but cannot send packets to other devices in the network.

**Note**

If PortFast is not enabled on the port, the port is forced to the bidirectional state.

When you configure a port as unidirectional by using the **authentication control-direction in** interface configuration command, the port changes to the spanning-tree forwarding state. The port can send packets to the host but cannot receive packets from the host.

When you configure a port as bidirectional by using the **authentication control-direction both** interface configuration command, the port is access-controlled in both directions. The port does not receive packets from or send packets to the host.

802.1x Authentication with MAC Authentication Bypass

You can configure the switch to authorize clients based on the client MAC address (see [Figure 1-2 on page 1-4](#)) by using the MAC authentication bypass feature. For example, you can enable this feature on 802.1x ports connected to devices such as printers.

If 802.1x authentication times out while waiting for an EAPOL response from the client, the switch tries to authorize the client by using MAC authentication bypass.

When the MAC authentication bypass feature is enabled on an 802.1x port, the switch uses the MAC address as the client identity. The authentication server has a database of client MAC addresses that are allowed network access. After detecting a client on an 802.1x port, the switch waits for an Ethernet

packet from the client. The switch sends the authentication server a RADIUS-access/request frame with a username and password based on the MAC address. If authorization succeeds, the switch grants the client access to the network. If authorization fails, the switch assigns the port to the guest VLAN if one is configured.

If an EAPOL packet is detected on the interface during the lifetime of the link, the switch determines that the device connected to that interface is an 802.1x-capable supplicant and uses 802.1x authentication (not MAC authentication bypass) to authorize the interface. EAPOL history is cleared if the interface link status goes down.

If the switch already authorized a port by using MAC authentication bypass and detects an 802.1x supplicant, the switch does not unauthorize the client connected to the port. When re-authentication occurs, the switch uses 802.1x authentication as the preferred re-authentication process if the previous session ended because the Termination-Action RADIUS attribute value is `DEFAULT`.

Clients that were authorized with MAC authentication bypass can be re-authenticated. The re-authentication process is the same as that for clients that were authenticated with 802.1x. During re-authentication, the port remains in the previously assigned VLAN. If re-authentication is successful, the switch keeps the port in the same VLAN. If re-authentication fails, the switch assigns the port to the guest VLAN, if one is configured.

If re-authentication is based on the Session-Timeout RADIUS attribute (Attribute[27]) and the Termination-Action RADIUS attribute (Attribute [29]) and if the Termination-Action RADIUS attribute (Attribute [29]) action is *Initialize*, (the attribute value is `DEFAULT`), the MAC authentication bypass session ends, and connectivity is lost during re-authentication. If MAC authentication bypass is enabled and the 802.1x authentication times out, the switch uses the MAC authentication bypass feature to initiate re-authorization. For more information about these AV pairs, see RFC 3580, “802.1X Remote Authentication Dial In User Service (RADIUS) Usage Guidelines.”

MAC authentication bypass interacts with the features:

- 802.1x authentication—You can enable MAC authentication bypass only if 802.1x authentication is enabled on the port.
- Guest VLAN—If a client has an invalid MAC address identity, the switch assigns the client to a guest VLAN if one is configured.
- Restricted VLAN—This feature is not supported when the client connected to an 802.1x port is authenticated with MAC authentication bypass.
- Port security—See the “[802.1x Authentication with Port Security](#)” section on page 1-27.
- Voice VLAN—See the “[802.1x Authentication with Voice VLAN Ports](#)” section on page 1-26.
- VLAN Membership Policy Server (VMPS)—802.1x and VMPS are mutually exclusive.
- Private VLAN—You can assign a client to a private VLAN.
- Network Edge Access Topology (NEAT)MAB and NEAT are mutually exclusive. You cannot enable MAB when NEAT is enabled on an interface, and you cannot enable NEAT when MAB is enabled on an interface.

For more configuration information, see the “[Authentication Manager](#)” section on page 1-7.

Cisco IOS Release 12.2(55)SE and later supports filtering of verbose MAB system messages. See the “[Authentication Manager CLI Commands](#)” section on page 1-9.

802.1x User Distribution

You can configure 802.1x user distribution to load-balance users with the same group name across multiple different VLANs.

The VLANs are either supplied by the RADIUS server or configured through the switch CLI under a VLAN group name.

- Configure the RADIUS server to send more than one VLAN name for a user. The multiple VLAN names can be sent as part of the response to the user. The 802.1x user distribution tracks all the users in a particular VLAN and achieves load balancing by moving the authorized user to the least populated VLAN.
- Configure the RADIUS server to send a VLAN group name for a user. The VLAN group name can be sent as part of the response to the user. You can search for the selected VLAN group name among the VLAN group names that you configured by using the switch CLI. If the VLAN group name is found, the corresponding VLANs under this VLAN group name are searched to find the least populated VLAN. Load balancing is achieved by moving the corresponding authorized user to that VLAN.



Note The RADIUS server can send the VLAN information in any combination of VLAN-IDs, VLAN names, or VLAN groups.

802.1x User Distribution Configuration Guidelines

- Confirm that at least one VLAN is mapped to the VLAN group.
- You can map more than one VLAN to a VLAN group.
- You can modify the VLAN group by adding or deleting a VLAN.
- When you clear an existing VLAN from the VLAN group name, none of the authenticated ports in the VLAN are cleared, but the mappings are removed from the existing VLAN group.
- If you clear the last VLAN from the VLAN group name, the VLAN group is cleared.
- You can clear a VLAN group even when the active VLANs are mapped to the group. When you clear a VLAN group, none of the ports or users that are in the authenticated state in any VLAN within the group are cleared, but the VLAN mappings to the VLAN group are cleared.

For more information, see the [“Configuring 802.1x User Distribution”](#) section on page 1-57.

Network Admission Control Layer 2 802.1x Validation



Note To use Network Admission Control, the switch must be running the LAN base image.

The switch supports the Network Admission Control (NAC) Layer 2 802.1x validation, which checks the antivirus condition or *posture* of endpoint systems or clients before granting the devices network access. With NAC Layer 2 802.1x validation, you can do these tasks:

- Download the Session-Timeout RADIUS attribute (Attribute[27]) and the Termination-Action RADIUS attribute (Attribute[29]) from the authentication server.

- Set the number of seconds between re-authentication attempts as the value of the Session-Timeout RADIUS attribute (Attribute[27]) and get an access policy against the client from the RADIUS server.
- Set the action to be taken when the switch tries to re-authenticate the client by using the Termination-Action RADIUS attribute (Attribute[29]). If the value is the *DEFAULT* or is not set, the session ends. If the value is RADIUS-Request, the re-authentication process starts.
- Set the list of VLAN number or name or VLAN group name as the value of the Tunnel Group Private ID (Attribute[81]) and the preference for the VLAN number or name or VLAN group name as the value of the Tunnel Preference (Attribute[83]). If you do not configure the Tunnel Preference, the first Tunnel Group Private ID (Attribute[81]) attribute is picked up from the list.
- View the NAC posture token, which shows the posture of the client, by using the **show authentication** or **show dot1x** privileged EXEC command.
- Configure secondary private VLANs as guest VLANs.

Configuring NAC Layer 2 802.1x validation is similar to configuring 802.1x port-based authentication except that you must configure a posture token on the RADIUS server. For information about configuring NAC Layer 2 802.1x validation, see the [“Configuring NAC Layer 2 802.1x Validation” section on page 1-58](#) and the [“Configuring Periodic Re-Authentication” section on page 1-45](#).

For more information about NAC, see the *Network Admission Control Software Configuration Guide*.

For more configuration information, see the [“Authentication Manager” section on page 1-7](#).

Flexible Authentication Ordering

You can use flexible authentication ordering to configure the order of methods that a port uses to authenticate a new host. MAC authentication bypass and 802.1x can be the primary or secondary authentication methods, and web authentication can be the fallback method if either or both of those authentication attempts fail. For more information see the [“Configuring Flexible Authentication Ordering” section on page 1-63](#).

Open1x Authentication

Open1x authentication allows a device to access a port before that device is authenticated. When open authentication is configured, a new host can pass traffic according to the access control list (ACL) defined on the port. After the host is authenticated, the policies configured on the RADIUS server are applied to that host.

You can configure open authentication with these scenarios:

- Single-host mode with open authentication—Only one user is allowed network access before and after authentication.
- MDA mode with open authentication—Only one user in the voice domain and one user in the data domain are allowed.
- Multiple-hosts mode with open authentication—Any host can access the network.
- Multiple-authentication mode with open authentication—Similar to MDA, except multiple hosts can be authenticated.

For more information see the [“Configuring the Host Mode” section on page 1-44](#).

**Note**

If open authentication is configured, it takes precedence over other authentication controls. This means that if you use the **authentication open** interface configuration command, the port will grant access to the host irrespective of the **authentication port-control** interface configuration command.

Using Voice Aware 802.1x Security

**Note**

To use voice aware IEEE 802.1x authentication, the switch must be running the LAN base image.

You use the voice aware 802.1x security feature to configure the switch to disable only the VLAN on which a security violation occurs, whether it is a data or voice VLAN. When an attempt to authenticate the data client caused a security violation in previous releases, the entire port shut down, resulting in a complete loss of connectivity.

You can use this feature where a PC is connected to the IP phone. A security violation found on the data VLAN shuts down only the data VLAN. The traffic on the voice VLAN continues without interruption.

For information on configuring voice aware 802.1x security, see the [“Configuring Voice Aware 802.1x Security” section on page 1-39](#).

802.1x Supplicant and Authenticator Switches with Network Edge Access Topology (NEAT)

The Network Edge Access Topology (NEAT) feature extends identity to areas outside the wiring closet (such as conference rooms). This allows any type of device to authenticate on the port.

- 802.1x switch supplicant: You can configure a switch to act as a supplicant to another switch by using the 802.1x supplicant feature. This configuration is helpful in a scenario, where, for example, a switch is outside a wiring closet and is connected to an upstream switch through a trunk port. A switch configured with the 802.1x switch supplicant feature authenticates with the upstream switch for secure connectivity.

Once the supplicant switch authenticates successfully the port mode changes from access to trunk.


- If the access VLAN is configured on the authenticator switch, it becomes the native VLAN for the trunk port after successful authentication.

You can enable MDA or multiauth mode on the authenticator switch interface that connects to one more supplicant switches. Multihost mode is not supported on the authenticator switch interface.

Use the **dot1x supplicant force-multicast** global configuration command on the supplicant switch for Network Edge Access Topology (NEAT) to work in all host modes.

- Host Authorization: Ensures that only traffic from authorized hosts (connecting to the switch with supplicant) is allowed on the network. The switches use Client Information Signalling Protocol (CISP) to send the MAC addresses connecting to the supplicant switch to the authenticator switch, as shown in [Figure 1-6](#).
- Auto enablement: Automatically enables trunk configuration on the authenticator switch, allowing user traffic from multiple VLANs coming from supplicant switches. Configure the `cisco-av-pair as device-traffic-class=switch` at the ACS. (You can configure this under the `group` or the `user` settings.)

Figure 1-6 Authenticator and Supplicant Switch using CISP



1	Workstations (clients)	2	Supplicant switch (outside wiring closet)
3	Authenticator switch	4	Access control server (ACS)
5	Trunk port		

Guidelines

- You can configure NEAT ports with the same configurations as the other authentication ports. When the supplicant switch authenticates, the port mode is changed from *access* to *trunk* based on the switch vendor-specific attributes (VSAs). (device-traffic-class=switch).
- The VSA changes the authenticator switch port mode from *access* to *trunk* and enables 802.1x trunk encapsulation and the access VLAN if any would be converted to a native trunk VLAN. VSA does not change any of the port configurations on the supplicant.
- To change the host mode *and* the apply a standard port configuration on the authenticator switch port, you can also use Auto Smartports user-defined macros, instead of the switch VSA. This allows you to remove unsupported configurations on the authenticator switch port and to change the port mode from *access* to *trunk*. For information, see the *AutoSmartports Configuration Guide*.

For more information, see the “[Configuring an Authenticator and a Supplicant Switch with NEAT](#)” section on page 1-59.

Using IEEE 802.1x Authentication with ACLs and the RADIUS Filter-Id Attribute



Note

To use IEEE 802.1x authentication with ACLs and the *Filter-Id* attribute, the switch must be running the LAN base image.

The switch supports both IP standard and IP extended port access control lists (ACLs) applied to ingress ports.

- ACLs that you configure
- ACLs from the Access Control Server (ACS)

An IEEE 802.1x port in single-host mode uses ACLs from the ACS to provide different levels of service to an IEEE 802.1x-authenticated user. When the RADIUS server authenticates this type of user and port, it sends ACL attributes based on the user identity to the switch. The switch applies the attributes to the port for the duration of the user session. If the session is over, authentication fails, or a link fails, the port becomes unauthorized, and the switch removes the ACL from the port.

Only IP standard and IP extended port ACLs from the ACS support the Filter-Id attribute. It specifies the name or number of an ACL. The Filter-id attribute can also specify the direction (inbound or outbound) and a user or a group to which the user belongs.

- The Filter-Id attribute for the user takes precedence over that for the group.
- If a Filter-Id attribute from the ACS specifies an ACL that is already configured, it takes precedence over a user-configured ACL.
- If the RADIUS server sends more than one Filter-Id attribute, only the last attribute is applied.

If the Filter-Id attribute is not defined on the switch, authentication fails, and the port returns to the unauthorized state.

Common Session ID

Authentication manager uses a single session ID (referred to as a common session ID) for a client no matter which authentication method is used. This ID is used for all reporting purposes, such as the show commands and MIBs. The session ID appears with all per-session syslog messages.

The session ID includes:

- The IP address of the Network Access Device (NAD)
- A monotonically increasing unique 32 bit integer
- The session start time stamp (a 32 bit integer)

This example shows how the session ID appears in the output of the **show authentication** command. The session ID in this example is 160000050000000B288508E5:

```
Switch# show authentication sessions
```

Interface	MAC Address	Method	Domain	Status	Session ID
Fa4/0/4	0000.0000.0203	mab	DATA	Authz Success	160000050000000B288508E5

This is an example of how the session ID appears in the syslog output. The session ID in this example is also 160000050000000B288508E5:

```
1w0d: %AUTHMGR-5-START: Starting 'mab' for client (0000.0000.0203) on Interface Fa4/0/4
AuditSessionID 160000050000000B288508E5
1w0d: %MAB-5-SUCCESS: Authentication successful for client (0000.0000.0203) on Interface
Fa4/0/4 AuditSessionID 160000050000000B288508E5
1w0d: %AUTHMGR-7-RESULT: Authentication result 'success' from 'mab' for client
(0000.0000.0203) on Interface Fa4/0/4 AuditSessionID 160000050000000B288508E5
```

The session ID is used by the NAD, the AAA server, and other report-analyzing applications to identify the client. The ID appears automatically. No configuration is required.

Configuring 802.1x Authentication

- [Default 802.1x Authentication Configuration, page 1-35](#)
- [802.1x Authentication Configuration Guidelines, page 1-36](#)
- [Configuring 802.1x Readiness Check, page 1-38 \(optional\)](#)
- [Configuring Voice Aware 802.1x Security, page 1-39 \(optional\)](#)
- [Configuring 802.1x Violation Modes, page 1-40 \(optional\)](#)
- [Configuring 802.1x Authentication, page 1-41 \(optional\)](#)
- [Configuring the Switch-to-RADIUS-Server Communication, page 1-43 \(required\)](#)
- [Configuring the Host Mode, page 1-44 \(optional\)](#)
- [Configuring Periodic Re-Authentication, page 1-45 \(optional\)](#)
- [Manually Re-Authenticating a Client Connected to a Port, page 1-46 \(optional\)](#)
- [Changing the Quiet Period, page 1-46 \(optional\)](#)
- [Changing the Switch-to-Client Retransmission Time, page 1-47 \(optional\)](#)
- [Setting the Switch-to-Client Frame-Retransmission Number, page 1-47 \(optional\)](#)
- [Setting the Re-Authentication Number, page 1-48 \(optional\)](#)
- [Configuring 802.1x Accounting, page 1-50 \(optional\)](#)
- [Enabling MAC Move, page 1-49 \(optional\)](#)
- [Enabling MAC Replace, page 1-49 \(optional\)](#)
- [Configuring a Guest VLAN, page 1-51 \(optional\)](#)
- [Configuring a Restricted VLAN, page 1-52 \(optional\)](#)
- [Configuring the Inaccessible Authentication Bypass Feature, page 1-54 \(optional\)](#)
- [Configuring 802.1x Authentication with Wake-on-LAN, page 1-56 \(optional\)](#)
- [Configuring MAC Authentication Bypass, page 1-56 \(optional\)](#)
- [Configuring NAC Layer 2 802.1x Validation, page 1-58 \(optional\)](#)
- [Configuring an Authenticator and a Supplicant Switch with NEAT, page 1-59 \(optional\)](#)
- [Configuring 802.1x Authentication with Downloadable ACLs and Redirect URLs, page 1-60 \(optional\)](#)
- [Configuring Flexible Authentication Ordering, page 1-63 \(optional\)](#)
- [Disabling 802.1x Authentication on the Port, page 1-65 \(optional\)](#)
- [Resetting the 802.1x Authentication Configuration to the Default Values, page 1-65 \(optional\)](#)

Default 802.1x Authentication Configuration

Table 1-4 Default 802.1x Authentication Configuration

Feature	Default Setting
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.
AAA	Disabled.
RADIUS server <ul style="list-style-type: none"> IP address UDP authentication port Key 	<ul style="list-style-type: none"> None specified. 1812. None specified.
Host mode	Single-host mode.
Control direction	Bidirectional control.
Periodic re-authentication	Disabled.
Number of seconds between re-authentication attempts	3600 seconds.
Re-authentication number	2 times (number of times that the switch restarts the authentication process before the port changes to the unauthorized state).
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).
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.) You can change this timeout period by using the authentication timer server interface configuration command.
Inactivity timeout	Disabled.
Guest VLAN	None specified.
Inaccessible authentication bypass	Disabled.
Restricted VLAN	None specified.
Authenticator (switch) mode	None specified.
MAC authentication bypass	Disabled.
Voice-aware security	Disabled

802.1x Authentication Configuration Guidelines

- [802.1x Authentication, page 1-36](#)
- [VLAN Assignment, Guest VLAN, Restricted VLAN, and Inaccessible Authentication Bypass, page 1-37](#)
- [MAC Authentication Bypass, page 1-37](#)
- [Maximum Number of Allowed Devices Per Port, page 1-38](#)

802.1x Authentication

- When IEEE 802.1x authentication is enabled, ports are authenticated before any other Layer 2 feature is enabled.
- If the VLAN to which an 802.1x-enabled port is assigned changes, this change is transparent and does not affect the switch. For example, this change occurs if a port is assigned to a RADIUS server-assigned VLAN and is then assigned to a different VLAN after re-authentication.

If the VLAN to which an 802.1x port is assigned to shut down, disabled, or removed, the port becomes unauthorized. For example, the port is unauthorized after the access VLAN to which a port is assigned shuts down or is removed.
- The IEEE 802.1x protocol is supported on Layer 2 static-access ports and voice VLAN ports, but it is not supported on these port types:
 - Dynamic ports—A port in dynamic mode can negotiate with its neighbor to become a trunk port. If you try to enable 802.1x authentication on a dynamic port, an error message appears, and 802.1x authentication 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 authentication on a dynamic-access (VLAN Query Protocol [VQP]) port, an error message appears, and 802.1x authentication 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 or a not-yet-active member of an EtherChannel as an 802.1x port. If you try to enable 802.1x authentication on an EtherChannel port, an error message appears, and 802.1x authentication is not enabled.
 - Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) destination ports—You can enable 802.1x authentication on a port that is a SPAN or RSPAN destination port. However, 802.1x authentication is disabled until the port is removed as a SPAN or RSPAN destination port. You can enable 802.1x authentication on a SPAN or RSPAN source port.
- Before globally enabling 802.1x authentication on a switch by entering the **dot1x system-auth-control** global configuration command, remove the EtherChannel configuration from the interfaces on which 802.1x authentication and EtherChannel are configured.
- Cisco IOS Release 12.2(55)SE and later supports filtering of system messages related to 802.1x authentication. See the [“Authentication Manager CLI Commands” section on page 1-9](#).

**Note**

The Catalyst 2960-S switch does not support Cisco Discovery Protocol (CDP) bypass.

VLAN Assignment, Guest VLAN, Restricted VLAN, and Inaccessible Authentication Bypass

- When 802.1x authentication is enabled on a port, you cannot configure a port VLAN that is equal to a voice VLAN.
- The 802.1x authentication with VLAN assignment feature is not supported on trunk ports, dynamic ports, or with dynamic-access port assignment through a VMPS.
- 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 trunk ports; it is supported only on access ports.
- After you configure a guest VLAN for an 802.1x port to which a DHCP client is connected, you might need to get a host IP address from a DHCP server. You can change the settings for restarting the 802.1x authentication process on the switch before the DHCP process on the client times out and tries to get a host IP address from the DHCP server. Decrease the settings for the 802.1x authentication process (**authentication timer inactivity** and **authentication timer reauthentication** interface configuration commands). The amount to decrease the settings depends on the connected 802.1x client type.
- When configuring the inaccessible authentication bypass feature, follow these guidelines:
 - The feature is supported on 802.1x port in single-host mode and multihosts mode.
 - If the client is running Windows XP and the port to which the client is connected is in the critical-authentication state, Windows XP might report that the interface is not authenticated.
 - If the Windows XP client is configured for DHCP and has an IP address from the DHCP server, receiving an EAP-Success message on a critical port might not re-initiate the DHCP configuration process.
 - You can configure the inaccessible authentication bypass feature and the restricted VLAN on an 802.1x port. If the switch tries to re-authenticate a critical port in a restricted VLAN and all the RADIUS servers are unavailable, switch changes the port state to the critical authentication state and remains in the restricted VLAN.
- You can configure any VLAN except an RSPAN VLAN or a voice VLAN as an 802.1x restricted VLAN. The restricted VLAN feature is not supported on trunk ports; it is supported only on access ports.

MAC Authentication Bypass

- Unless otherwise stated, the MAC authentication bypass guidelines are the same as the 802.1x authentication guidelines. For more information, see the [“802.1x Authentication” section on page 1-36](#).
- If you disable MAC authentication bypass from a port after the port has been authorized with its MAC address, the port state is not affected.
- If the port is in the unauthorized state and the client MAC address is not the authentication-server database, the port remains in the unauthorized state. However, if the client MAC address is added to the database, the switch can use MAC authentication bypass to re-authorize the port.
- If the port is in the authorized state, the port remains in this state until re-authorization occurs.
- You can configure a timeout period for hosts that are connected by MAC authentication bypass but are inactive. The range is 1 to 65535 seconds.

Maximum Number of Allowed Devices Per Port

This is the maximum number of devices allowed on an 802.1x-enabled port:

- In single-host mode, only one device is allowed on the access VLAN. If the port is also configured with a voice VLAN, an unlimited number of Cisco IP phones can send and receive traffic through the voice VLAN.
- In multidomain authentication (MDA) mode, one device is allowed for the access VLAN, and one IP phone is allowed for the voice VLAN.
- In multiple-host mode, only one 802.1x supplicant is allowed on the port, but an unlimited number of non-802.1x hosts are allowed on the access VLAN. An unlimited number of devices are allowed on the voice VLAN.

Configuring 802.1x Readiness Check

The 802.1x readiness check monitors 802.1x activity on all the switch ports and displays information about the devices connected to the ports that support 802.1x. You can use this feature to determine if the devices connected to the switch ports are 802.1x-capable.

The 802.1x readiness check is allowed on all ports that can be configured for 802.1x. The readiness check is not available on a port that is configured as **dot1x force-unauthorized**.

Follow these guidelines to enable the readiness check on the switch:

- The readiness check is typically used before 802.1x is enabled on the switch.
- If you use the **dot1x test eapol-capable** privileged EXEC command without specifying an interface, all the ports on the switch stack are tested.
- When you configure the **dot1x test eapol-capable** command on an 802.1x-enabled port, and the link comes up, the port queries the connected client about its 802.1x capability. When the client responds with a notification packet, it is 802.1x-capable. A syslog message is generated if the client responds within the timeout period. If the client does not respond to the query, the client is not 802.1x-capable. No syslog message is generated.
- The readiness check can be sent on a port that handles multiple hosts (for example, a PC that is connected to an IP phone). A syslog message is generated for each of the clients that respond to the readiness check within the timer period.

Beginning in privileged EXEC mode, follow these steps to enable the 802.1x readiness check on the switch:

	Command	Purpose
Step 1	dot1x test eapol-capable [interface <i>interface-id</i>]	Enable the 802.1x readiness check on the switch. (Optional) For <i>interface-id</i> specify the port on which to check for 802.1x readiness. Note If you omit the optional interface keyword, all interfaces on the switch are tested.
Step 1	configure terminal	(Optional) Enter global configuration mode.
Step 2	dot1x test timeout <i>timeout</i>	(Optional) Configure the timeout used to wait for EAPOL response. The range is from 1 to 65535 seconds. The default is 10 seconds.

	Command	Purpose
Step 3	end	(Optional) Return to privileged EXEC mode.
Step 4	show running-config	(Optional) Verify your modified timeout values.

This example shows how to enable a readiness check on a switch to query a port. It also shows the response received from the queried port verifying that the device connected to it is 802.1x-capable:

```
Switch# dot1x test eapol-capable interface gigabitethernet1/0/13
Switch# dot1x test eapol-capable interface gigabitethernet0/13

DOT1X_PORT_EAPOL_CAPABLE:DOT1X: MAC 00-01-02-4b-f1-a3 on gigabitethernet1/0/13 is EAPOL
capable
DOT1X_PORT_EAPOL_CAPABLE:DOT1X: MAC 00-01-02-4b-f1-a3 on gigabitethernet0/13 is EAPOL
capable
```

Configuring Voice Aware 802.1x Security



Note

To use voice aware IEEE 802.1x authentication, the switch must be running the LAN base image.

You use the voice aware 802.1x security feature on the switch to disable only the VLAN on which a security violation occurs, whether it is a data or voice VLAN. You can use this feature in IP phone deployments where a PC is connected to the IP phone. A security violation found on the data VLAN results in the shutdown of only the data VLAN. The traffic on the voice VLAN flows through the switch without interruption.

Follow these guidelines to configure voice aware 802.1x voice security on the switch:

- You enable voice aware 802.1x security by entering the **errdisable detect cause security-violation shutdown vlan** global configuration command. You disable voice aware 802.1x security by entering the **no** version of this command. This command applies to all 802.1x-configured ports in the switch.



Note

If you do not include the **shutdown vlan** keywords, the entire port is shut down when it enters the error-disabled state.

- If you use the **errdisable recovery cause security-violation** global configuration command to configure error-disabled recovery, the port is automatically re-enabled. If error-disabled recovery is not configured for the port, you re-enable it by using the **shutdown** and **no-shutdown** interface configuration commands.
- You can re-enable individual VLANs by using the **clear errdisable interface interface-id vlan [vlan-list]** privileged EXEC command. If you do not specify a range, all VLANs on the port are enabled.

Beginning in privileged EXEC mode, follow these steps to enable voice aware 802.1x security:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	errdisable detect cause security-violation shutdown vlan	Shut down any VLAN on which a security violation error occurs. Note If the shutdown vlan keywords are not included, the entire port enters the error-disabled state and shuts down.
Step 3	errdisable recovery cause security-violation	(Optional) Enable automatic per-VLAN error recovery.
Step 4	clear errdisable interface <i>interface-id</i> vlan [<i>vlan-list</i>]	(Optional) Reenable individual VLANs that have been error disabled. <ul style="list-style-type: none"> For <i>interface-id</i> specify the port on which to reenable individual VLANs. (Optional) For <i>vlan-list</i> specify a list of VLANs to be re-enabled. If <i>vlan-list</i> is not specified, all VLANs are re-enabled.
Step 5	shutdown no-shutdown	(Optional) Re-enable an error-disabled VLAN, and clear all error-disable indications.
Step 6	end	Return to privileged EXEC mode.
Step 7	show errdisable detect	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to configure the switch to shut down any VLAN on which a security violation error occurs:

```
Switch(config)# errdisable detect cause security-violation shutdown vlan
```

This example shows how to re-enable all VLANs that were error disabled on port Gigabit Ethernet 0/2.

```
Switch# clear errdisable interface gigabitethernet0/2 vlan
```

You can verify your settings by entering the **show errdisable detect** privileged EXEC command.

Configuring 802.1x Violation Modes



Note

To configure violation modes, the switch must be running the LAN base image.

You can configure an 802.1x port so that it shuts down, generates a syslog error, or discards packets from a new device when:

- a device connects to an 802.1x-enabled port
- the maximum number of allowed about devices have been authenticated on the port

Beginning in privileged EXEC mode, follow these steps to configure the security violation actions on the switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa new-model	Enable AAA.
Step 3	aaa authentication dot1x {default} <i>method1</i>	<p>Create an 802.1x authentication method list.</p> <p>To create a default list to use when a named list is <i>not</i> specified in the authentication command, use the default keyword followed by the method that is to be used in default situations. The default method list is automatically applied to all ports.</p> <p>For <i>method1</i>, enter the group radius keywords to use the list of all RADIUS servers for authentication.</p> <p>Note Though other keywords are visible in the command-line help string, only the group radius keywords are supported.</p>
Step 4	interface <i>interface-id</i>	Specify the port connected to the client that is to be enabled for 802.1x authentication, and enter interface configuration mode.
Step 5	switchport mode access	Set the port to access mode.
Step 6	authentication violation {shutdown restrict protect replace}	<p>Configure the violation mode. The keywords have these meanings:</p> <ul style="list-style-type: none"> • shutdown—Error disable the port. • restrict—Generate a syslog error. • protect—Drop packets from any new device that sends traffic to the port. • replace—Removes the current session and authenticates with the new host.
Step 7	end	Return to privileged EXEC mode.
Step 8	show authentication	Verify your entries.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Configuring 802.1x Authentication

To configure 802.1x port-based authentication, you must enable authentication, authorization, and accounting (AAA) and specify the authentication method list. A method list describes the sequence and authentication method to be queried to authenticate a user.

To allow VLAN assignment, you must enable AAA authorization to configure the switch for all network-related service requests.

This is the 802.1x AAA process:

-
- Step 1 A user connects to a port on the switch.
 - Step 2 Authentication is performed.
 - Step 3 VLAN assignment is enabled, as appropriate, based on the RADIUS server configuration.
 - Step 4 The switch sends a start message to an accounting server.

- Step 5** Re-authentication is performed, as necessary.
- Step 6** The switch sends an interim accounting update to the accounting server, that is based on the result of re-authentication.
- Step 7** The user disconnects from the port.
- Step 8** The switch sends a stop message to the accounting server.

Beginning in privileged EXEC mode, follow these steps to configure 802.1x port-based authentication:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa new-model	Enable AAA.
Step 3	aaa authentication dot1x {default} method1	<p>Create an 802.1x authentication method list.</p> <p>To create a default list to use when a named list is <i>not</i> specified in the authentication command, use the default keyword followed by the method to use in default situations. The default method list is automatically applied to all ports.</p> <p>For <i>method1</i>, enter the group radius keywords to use the list of all RADIUS servers for authentication.</p> <p>Note Though other keywords are visible in the command-line help string, only the group radius keywords are supported.</p>
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 to use user-RADIUS authorization for all network-related service requests, such as VLAN assignment.
Step 6	radius-server host ip-address	(Optional) Specify the IP address of the RADIUS server.
Step 7	radius-server key string	(Optional) Specify the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server.
Step 8	interface interface-id	Specify the port connected to the client to enable for 802.1x authentication, and enter interface configuration mode.
Step 9	switchport mode access	(Optional) Set the port to access mode only if you configured the RADIUS server in Step 6 and Step 7.
Step 10	authentication port-control auto	<p>Enable 802.1x authentication on the port.</p> <p>For feature interaction information, see the “802.1x Authentication Configuration Guidelines” section on page 1-36.</p>
Step 11	dot1x pae authenticator	Set the interface Port Access Entity to act only as an authenticator and ignore messages meant for a supplicant.
Step 12	end	Return to privileged EXEC mode.
Step 13	show authentication	Verify your entries.
Step 14	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Configuring the Switch-to-RADIUS-Server Communication

RADIUS security servers are identified by their hostname or IP address, hostname and specific UDP port numbers, or IP address and specific UDP port numbers. The combination of the IP address and the 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 in which 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 { <i>hostname</i> <i>ip-address</i> } auth-port <i>port-number</i> key <i>string</i>	<p>Configure the RADIUS server parameters.</p> <p>For <i>hostname</i> <i>ip-address</i>, specify the hostname or IP address of the remote RADIUS server.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>If you want to use multiple RADIUS servers, re-enter this command.</p>
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 clear 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 1-36.

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, see the RADIUS server documentation.

Configuring the Host Mode

Beginning in privileged EXEC mode, follow these steps to allow a single host (client) or multiple hosts on an 802.1x-authorized port that has the **authentication port-control** interface configuration command set to **auto**. Use the **multi-domain** keyword to configure multidomain authentication (MDA) to enable authentication of both a host and a voice device, such as an IP phone (Cisco or non-Cisco) on the same switch port.

This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	radius-server vsa send authentication	Configure the network access server to recognize and use vendor-specific attributes (VSAs).
Step 3	interface <i>interface-id</i>	Specify the port to which multiple hosts are indirectly attached, and enter interface configuration mode.
Step 4	authentication host-mode [multi-auth multi-domain multi-host single-host]	<p>The keywords have these meanings:</p> <ul style="list-style-type: none"> multi-auth—Allow one client on the voice VLAN and multiple authenticated clients on the data VLAN. Each host is individually authenticated. <p>Note The multi-auth keyword is only available with the authentication host-mode command.</p> <ul style="list-style-type: none"> multi-host—Allow multiple hosts on an 802.1x-authorized port after a single host has been authenticated. multi-domain—Allow both a host and a voice device, such as an IP phone (Cisco or non-Cisco), to be authenticated on an 802.1x-authorized port. <p>Note You must configure the voice VLAN for the IP phone when the host mode is set to multi-domain. For more information, see Chapter 1, “Configuring Voice VLAN.”</p> <ul style="list-style-type: none"> single-host—Allow a single host (client) on an 802.1x-authorized port. <p>Make sure that the authentication port-control interface configuration command set is set to auto for the specified interface.</p>
Step 5	switchport voice vlan <i>vlan-id</i>	(Optional) Configure the voice VLAN.
Step 6	end	Return to privileged EXEC mode.
Step 7	show authentication 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 multiple hosts on the port, use the **no authentication host-mode** interface configuration command.

This example shows how to enable 802.1x authentication and to allow multiple hosts:

```
Switch(config)# interface gigabitethernet2/0/1
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# authentication port-control auto
Switch(config-if)# authentication host-mode multi-host
Switch(config-if)# end
```

This example shows how to enable MDA and to allow both a host and a voice device on the port:

```
Switch(config)# interface gigabitethernet2/0/1
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# authentication port-control auto
Switch(config-if)# authentication host-mode multi-domain
Switch(config-if)# switchport voice vlan 101
Switch(config-if)# end
```

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 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	authentication periodic	Enable periodic re-authentication of the client, which is disabled by default. Note The default value is 3600 seconds. To change the value of the reauthentication timer or to have the switch use a RADIUS-provided session timeout, enter the authentication timer reauthenticate command.
Step 4	authentication timer {{{inactivity reauthenticate}} {restart <i>value</i> }}	Set the number of seconds between re-authentication attempts. The authentication timer keywords have these meanings: <ul style="list-style-type: none"> • inactivity—Interval in seconds after which if there is no activity from the client then it is unauthorized • reauthenticate—Time in seconds after which an automatic re-authentication attempt is initiated • restart <i>value</i>—Interval in seconds after which an attempt is made to authenticate an unauthorized port This command affects the behavior of the switch only if periodic re-authentication is enabled.
Step 5	end	Return to privileged EXEC mode.

	Command	Purpose
Step 6	show authentication 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 authentication periodic** interface configuration command. To return to the default number of seconds between re-authentication attempts, use the **no authentication timer** 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)# authentication periodic
Switch(config-if)# authentication timer reauthenticate 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 1-45](#).

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

```
Switch# dot1x re-authenticate interface gigabitethernet2/0/1
Switch# dot1x re-authenticate interface gigabitethernet0/1
```

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 **authentication timer inactivity** interface configuration command controls the idle period. A failed client authentication might occur because the client provided an invalid password. You can provide a faster response time to the user by entering a number smaller 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	authentication timer inactivity <i>seconds</i>	Set the number of seconds that the switch remains in the quiet state after 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 authentication 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 authentication timer inactivity** interface configuration command.

This example shows how to set the quiet time on the switch to 30 seconds:

```
Switch(config-if)# authentication timer inactivity 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 configure, and enter interface configuration mode.
Step 3	authentication timer reauthenticate <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 5.
Step 4	end	Return to privileged EXEC mode.
Step 5	show authentication 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 time, use the **no authentication timer reauthenticate** 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)# authentication timer reauthenticate 60
```

Setting the Switch-to-Client Frame-Retransmission Number

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-reauth-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 authentication 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-reauth-req 5
```

Setting the Re-Authentication Number

You can also change the number of times that the switch restarts the authentication process before the port changes to the unauthorized state.



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 re-authentication 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 restarts the authentication process before the port changes to the unauthorized state. The range is 0 to 10; the default is 2.
Step 4	end	Return to privileged EXEC mode.
Step 5	show authentication 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 re-authentication number, use the **no dot1x max-reauth-req** interface configuration command.

This example shows how to set 4 as the number of times that the switch restarts the authentication process before the port changes to the unauthorized state:

```
Switch(config-if)# dot1x max-reauth-req 4
```

Enabling MAC Move

MAC move allows an authenticated host to move from one port on the switch to another.

Beginning in privileged EXEC mode, follow these steps to globally enable MAC move on the switch. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	authentication mac-move permit	Enable MAC move on the switch.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	(Optional) Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to globally enable MAC move on a switch:

```
Switch(config)# authentication mac-move permit
```

Enabling MAC Replace



Note

To enable MAC replace, the switch must be running the LAN base image.

MAC replace allows a host to replace an authenticated host on a port.

Beginning in privileged EXEC mode, follow these steps to enable MAC replace 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 port to be configured, and enter interface configuration mode.

	Command	Purpose
Step 3	authentication violation {protect replace restrict shutdown}	Use the replace keyword to enable MAC replace on the interface. The port removes the current session and initiates authentication with the new host. The other keywords have these effects: <ul style="list-style-type: none"> • protect: the port drops packets with unexpected MAC addresses without generating a system message. • restrict: violating packets are dropped by the CPU and a system message is generated. • shutdown: the port is error disabled when it receives an unexpected MAC address.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Saves your entries in the configuration file.

This example shows how to enable MAC replace on an interface:

```
Switch(config)# interface gigabitethernet2/0/2
Switch(config-if)# authentication violation replace
```

Configuring 802.1x Accounting

Enabling AAA system accounting with 802.1x accounting allows system reload events to be sent to the accounting RADIUS server for logging. The server can then infer that all active 802.1x sessions are closed.

Because RADIUS uses the unreliable UDP transport protocol, accounting messages might be lost due to poor network conditions. If the switch does not receive the accounting response message from the RADIUS server after a configurable number of retransmissions of an accounting request, this system message appears:

```
Accounting message %s for session %s failed to receive Accounting Response.
```

When the stop message is not sent successfully, this message appears:

```
00:09:55: %RADIUS-4-RADIUS_DEAD: RADIUS server 172.20.246.201:1645,1646 is not responding.
```



Note

You must configure the RADIUS server to perform accounting tasks, such as logging start, stop, and interim-update messages and time stamps. To turn on these functions, enable logging of “Update/Watchdog packets from this AAA client” in your RADIUS server Network Configuration tab. Next, enable “CVS RADIUS Accounting” in your RADIUS server System Configuration tab.

Beginning in privileged EXEC mode, follow these steps to configure 802.1x accounting after AAA is enabled on your switch. 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.

	Command	Purpose
Step 3	aaa accounting dot1x default start-stop group radius	Enable 802.1x accounting using the list of all RADIUS servers.
Step 4	aaa accounting system default start-stop group radius	(Optional) Enables system accounting (using the list of all RADIUS servers) and generates system accounting reload event messages when the switch reloads.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Saves your entries in the configuration file.

Use the **show radius statistics** privileged EXEC command to display the number of RADIUS messages that do not receive the accounting response message.

This example shows how to configure 802.1x accounting. The first command configures the RADIUS server, specifying 1813 as the UDP port for accounting:

```
Switch(config)# radius-server host 172.120.39.46 auth-port 1812 acct-port 1813 key rad123
Switch(config)# aaa accounting dot1x default start-stop group radius
Switch(config)# aaa accounting system default start-stop group radius
```

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 EAP request/identity frame. Clients that are 802.1x-capable but that fail authentication are not granted network access. 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 Authentication Configuration Guidelines ” section on page 1-36.
Step 3	switchport mode access	Set the port to access mode.
Step 4	authentication port-control auto	Enable 802.1x authentication on the port.
Step 5	authentication event no-response action authorize 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 RSPAN VLAN or a voice VLAN as an 802.1x guest VLAN.
Step 6	end	Return to privileged EXEC mode.
Step 7	show authentication 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 and remove the guest VLAN, use the **no authentication event no-response action authorize vlan *vlan-id*** interface configuration command. The port returns to the unauthorized state.

This example shows how to enable VLAN 2 as an 802.1x guest VLAN:

```
Switch(config)# interface gigabitethernet2/0/2
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# authentication event no-response action authorize vlan 2
```

This example shows how to set 3 as the quiet time on the switch, to set 15 as the number of seconds that the switch waits for a response to an EAP-request/identity frame from the client before re-sending the request, and to enable VLAN 2 as an 802.1x guest VLAN when an 802.1x port is connected to a DHCP client:

```
Switch(config-if)# authentication timer inactivity 3
Switch(config-if)# authentication timer reauthenticate 15
Switch(config-if)# authentication event no-response action authorize vlan 2
```

Configuring a Restricted VLAN

When you configure a restricted VLAN on a switch stack or a switch, clients that are 802.1x-compliant are moved into the restricted VLAN when the authentication server does not receive a valid username and password. The switch supports restricted VLANs only in single-host mode.

Beginning in privileged EXEC mode, follow these steps to configure a restricted 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 Authentication Configuration Guidelines” section on page 1-36.
Step 3	switchport mode access	Set the port to access mode.
Step 4	authentication port-control auto	Enable 802.1x authentication on the port.
Step 5	authentication event fail action authorize <i>vlan-id</i>	Specify an active VLAN as an 802.1x restricted VLAN. The range is 1 to 4094. You can configure any active VLAN except an RSPAN VLAN or a voice VLAN as an 802.1x restricted VLAN.
Step 6	end	Return to privileged EXEC mode.
Step 7	show authentication interface <i>interface-id</i>	(Optional) Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable and remove the restricted VLAN, use the **no authentication event fail action authorize *vlan-id*** interface configuration command. The port returns to the unauthorized state.

This example shows how to enable VLAN 2 as an 802.1x restricted VLAN:

```
Switch(config-if)# authentication event fail action authorize 2
```

You can configure the maximum number of authentication attempts allowed before a user is assigned to the restricted VLAN by using the **authentication event retry** *retry count* interface configuration command. The range of allowable authentication attempts is 1 to 3. The default is 3 attempts.

Beginning in privileged EXEC mode, follow these steps to configure the maximum number of allowed 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. For the supported port types, see the “ 802.1x Authentication Configuration Guidelines ” section on page 1-36.
Step 3	switchport mode access	Set the port to access mode.
Step 4	authentication port-control auto	Enable 802.1x authentication on the port.
Step 5	authentication event fail action authorize <i>vlan-id</i>	Specify an active VLAN as an 802.1x restricted VLAN. The range is 1 to 4094. You can configure any active VLAN except an RSPAN VLAN or a voice VLAN as an 802.1x restricted VLAN.
Step 6	authentication event retry <i>retry count</i>	Specify a number of authentication attempts to allow before a port moves to the restricted VLAN. The range is 1 to 3, and the default is 3.
Step 7	end	Return to privileged EXEC mode.
Step 8	show authentication interface <i>interface-id</i>	(Optional) Verify your entries.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default value, use the **no authentication event retry** interface configuration command.

This example shows how to set 2 as the number of authentication attempts allowed before the port moves to the restricted VLAN:

```
Switch(config-if)# authentication event retry 2
```

Configuring the Inaccessible Authentication Bypass Feature

You can configure the inaccessible bypass feature, also referred to as critical authentication or the AAA fail policy.

Beginning in privileged EXEC mode, follow these steps to configure the port as a critical port and enable the inaccessible authentication bypass feature. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	radius-server dead-criteria <i>time time tries tries</i>	<p>(Optional) Set the conditions that are used to decide when a RADIUS server is considered unavailable or <i>dead</i>.</p> <p>The range for <i>time</i> is from 1 to 120 seconds. The switch dynamically determines the default <i>seconds</i> value that is 10 to 60 seconds.</p> <p>The range for <i>tries</i> is from 1 to 100. The switch dynamically determines the default <i>tries</i> parameter that is 10 to 100.</p>
Step 3	radius-server deadtime <i>minutes</i>	(Optional) Set the number of minutes that a RADIUS server is not sent requests. The range is from 0 to 1440 minutes (24 hours). The default is 0 minutes.
Step 4	radius-server host <i>ip-address</i> [acct-port <i>udp-port</i>] [auth-port <i>udp-port</i>] [test username <i>name</i>] [idle-time <i>time</i>] [ignore-acct-port] [ignore-auth-port] [key <i>string</i>]	<p>(Optional) Configure the RADIUS server parameters by using these keywords:</p> <ul style="list-style-type: none"> • acct-port <i>udp-port</i>—Specify the UDP port for the RADIUS accounting server. The range for the UDP port number is from 0 to 65536. The default is 1646. • auth-port <i>udp-port</i>—Specify the UDP port for the RADIUS authentication server. The range for the UDP port number is from 0 to 65536. The default is 1645. <p>Note You should configure the UDP port for the RADIUS accounting server and the UDP port for the RADIUS authentication server to nondefault values.</p> <ul style="list-style-type: none"> • test username <i>name</i>—Enable automated testing of the RADIUS server status, and specify the username to be used. • idle-time <i>time</i>—Set the interval of time in minutes after which the switch sends test packets to the server. The range is from 1 to 35791 minutes. The default is 60 minutes (1 hour). • ignore-acct-port—Disable testing on the RADIUS-server accounting port. • ignore-auth-port—Disable testing on the RADIUS-server authentication port. • key <i>string</i>—Specify the authentication and encryption key for all RADIUS communication between the switch and the RADIUS daemon. <p>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.</p> <p>You can also configure the authentication and encryption key by using the radius-server key {0 <i>string</i> 7 <i>string</i> <i>string</i>} global configuration command.</p>

	Command	Purpose
Step 5	dot1x critical {eapol recovery delay <i>milliseconds</i>}	(Optional) Configure the parameters for inaccessible authentication bypass: eapol —Specify that the switch sends an EAPOL-Success message when the switch successfully authenticates the critical port. recovery delay <i>milliseconds</i> —Set the recovery delay period during which the switch waits to re-initialize a critical port when a RADIUS server that was unavailable becomes available. The range is from 1 to 10000 milliseconds. The default is 1000 milliseconds (a port can be re-initialized every second).
Step 6	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 Authentication Configuration Guidelines ” section on page 1-36.
Step 7	authentication event server dead action [authorize reinitialize] vlan <i>vlan-id</i>	Use these keywords to move hosts on the port if the RADIUS server is unreachable: <ul style="list-style-type: none"> authorize—Move any new hosts trying to authenticate to the user-specified critical VLAN. reinitialize—Move all authorized hosts on the port to the user-specified critical VLAN.
Step 8	authentication event server dead action {authorize reinitialize} vlan <i>vlan-id</i>	Enable the inaccessible authentication bypass feature, and use these keywords to configure the feature: <ul style="list-style-type: none"> authorize—Authorize the port. reinitialize—Reinitialize all authorized clients.
Step 9	authentication server dead action authorize [vlan]	Authorize the switch in access VLAN or configured VLAN (if the VLAN is specified) when the ACS server is down.
Step 10	end	Return to privileged EXEC mode.
Step 11	show authentication interface <i>interface-id</i>	(Optional) Verify your entries.
Step 12	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the RADIUS server default settings, use the **no radius-server dead-criteria**, the **no radius-server deadtime**, and the **no radius-server host** global configuration commands. To return to the default settings of inaccessible authentication bypass, use the **no authentication event server dead action {authorize | reinitialize}** global configuration command. To disable inaccessible authentication bypass, use the **no authentication event server dead action** interface configuration command.

This example shows how to configure the inaccessible authentication bypass feature:

```
Switch(config)# radius-server dead-criteria time 30 tries 20
Switch(config)# radius-server deadtime 60
Switch(config)# radius-server host 1.1.1.2 acct-port 1550 auth-port 1560 test username
user1 idle-time 30 key abc1234
Switch(config)# dot1x critical eapol
Switch(config)# dot1x critical recovery delay 2000
Switch(config)# interface gigabitethernet1/0/2
Switch(config)# radius-server deadtime 60
Switch(config-if)# dot1x critical
Switch(config-if)# dot1x critical recovery action reinitialize?
Switch(config-if)# dot1x critical vlan 20
Switch(config-if)# end
```

Configuring 802.1x Authentication with Wake-on-LAN

Beginning in privileged EXEC mode, follow these steps to enable 802.1x authentication with WoL. 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 Authentication Configuration Guidelines ” section on page 1-36.
Step 3	authentication control-direction { both in }	Enable 802.1x authentication with WoL on the port, and use these keywords to configure the port as bidirectional or unidirectional. <ul style="list-style-type: none"> • both—Sets the port as bidirectional. The port cannot receive packets from or send packets to the host. By default, the port is bidirectional. • in—Sets the port as unidirectional. The port can send packets to the host but cannot receive packets from the host.
Step 4	end	Return to privileged EXEC mode.
Step 5	show authentication 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 802.1x authentication with WoL, use the **no authentication control-direction** interface configuration command.

These examples show how to enable 802.1x authentication with WoL and set the port as bidirectional:

```
Switch(config-if)# authentication control-direction both
```

Configuring MAC Authentication Bypass

Beginning in privileged EXEC mode, follow these steps to enable MAC authentication bypass. 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 Authentication Configuration Guidelines ” section on page 1-36.
Step 3	authentication port-control auto	Enable 802.1x authentication on the port.
Step 4	authentication order [mab] { webauth }	Set the order of authentication methods. <ul style="list-style-type: none"> • mab—Add MAC authentication bypass (MAB) to the order of authentication methods. • webauth—Add web authentication to the order of authentication methods.
Step 5	end	Return to privileged EXEC mode.

	Command	Purpose
Step 6	show authentication 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 MAC authentication bypass, use the **no authentication order** interface configuration command.

This example shows how to enable MAC authentication bypass:

```
Switch(config-if)# authentication order
```

Configuring 802.1x User Distribution

Beginning in global configuration, follow these steps to configure a VLAN group and to map a VLAN to it:

	Command	Purpose
Step 1	vlan group <i>vlan-group-name</i> vlan-list <i>vlan-list</i>	Configure a VLAN group, and map a single VLAN or a range of VLANs to it.
Step 2	show vlan group all <i>vlan-group-name</i>	Verify the configuration.
Step 3	no vlan group <i>vlan-group-name</i> vlan-list <i>vlan-list</i>	Clear the VLAN group configuration or elements of the VLAN group configuration.

This example shows how to configure the VLAN groups, to map the VLANs to the groups, to and verify the VLAN group configurations and mapping to the specified VLANs:

```
Switch(config)# vlan group eng-dept vlan-list 10

switch(config)# show vlan group group-name eng-dept
Group Name                Vlans Mapped
-----
eng-dept                  10
switch# show dot1x vlan-group all
Group Name                Vlans Mapped
-----
eng-dept                  10
hr-dept                   20
```

This example shows how to add a VLAN to an existing VLAN group and to verify that the VLAN was added:

```
switch(config)# vlan group eng-dept vlan-list 30
switch(config)# show vlan group eng-dept
Group Name                Vlans Mapped
-----
eng-dept                  10, 30
```

This example shows how to remove a VLAN from a VLAN group:

```
switch# no vlan group eng-dept vlan-list 10
```

This example shows that when all the VLANs are cleared from a VLAN group, the VLAN group is cleared:

```
switch(config)# no vlan group eng-dept vlan-list 30
Vlan 30 is successfully cleared from vlan group eng-dept.
```

```
switch(config)# show vlan group group-name eng-dept
```

This example shows how to clear all the VLAN groups:

```
switch(config)# no vlan group eng-dept vlan-list all
switch(config)# show vlan-group all
```

For more information about these commands, see the *Cisco IOS Security Command Reference*.

Configuring NAC Layer 2 802.1x Validation

You can configure NAC Layer 2 802.1x validation, which is also referred to as 802.1x authentication with a RADIUS server.

Beginning in privileged EXEC mode, follow these steps to configure NAC Layer 2 802.1x validation. The 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	authentication event no-response action authorize 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 RSPAN VLAN, or a voice VLAN as an 802.1x guest VLAN.
Step 4	authentication periodic	Enable periodic re-authentication of the client, which is disabled by default.
Step 5	authentication timer reauthenticate	Set re-authentication attempt for the client (set to one hour). This command affects the behavior of the switch only if periodic re-authentication is enabled.
Step 6	end	Return to privileged EXEC mode.
Step 7	show authentication interface <i>interface-id</i>	Verify your 802.1x authentication configuration.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to configure NAC Layer 2 802.1x validation:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# authentication periodic
Switch(config-if)# authentication timer reauthenticate
```

Configuring an Authenticator and a Supplicant Switch with NEAT

Configuring this feature requires that one switch outside a wiring closet is configured as a supplicant and is connected to an authenticator switch.

For overview information, see the “[802.1x Supplicant and Authenticator Switches with Network Edge Access Topology \(NEAT\)](#)” section on page 1-31.



Note The *cisco-av-pairs* must be configured as *device-traffic-class=switch* on the ACS, which sets the interface as a trunk after the supplicant is successfully authenticated.

Beginning in privileged EXEC mode, follow these steps to configure a switch as an authenticator:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	cisp enable	Enable CISP.
Step 3	interface <i>interface-id</i>	Specify the port to be configured, and enter interface configuration mode.
Step 4	switchport mode access	Set the port mode to access .
Step 5	authentication port-control auto	Set the port-authentication mode to auto.
Step 6	dot1x pae authenticator	Configure the interface as a port access entity (PAE) authenticator.
Step 7	spanning-tree portfast	Enable Port Fast on an access port connected to a single workstation or server.
Step 8	end	Return to privileged EXEC mode.
Step 9	show running-config interface <i>interface-id</i>	Verify your configuration.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to configure a switch as an 802.1x authenticator:

```
Switch# configure terminal
Switch(config)# cisp enable
Switch(config)# interface gigabitethernet2/0/1
Switch(config-if)# switchport mode access
Switch(config-if)# authentication port-control auto
Switch(config-if)# dot1x pae authenticator
Switch(config-if)# spanning-tree portfast trunk
```

Beginning in privileged EXEC mode, follow these steps to configure a switch as a supplicant:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	cisp enable	Enable CISP.
Step 3	dot1x credentials <i>profile</i>	Create 802.1x credentials profile. This must be attached to the port that is configured as supplicant.
Step 4	username <i>suppswitch</i>	Create a username.

	Command	Purpose
Step 5	<code>password <i>password</i></code>	Create a password for the new username.
Step 6	<code>dot1x supplicant force-multicast</code>	Force the switch to send <i>only</i> multicast EAPOL packets when it receives either unicast or multicast packets. This also allows NEAT to work on the supplicant switch in all host modes.
Step 7	<code>interface <i>interface-id</i></code>	Specify the port to be configured, and enter interface configuration mode.
Step 8	<code>switchport trunk encapsulation dot1q</code>	Set the port to trunk mode.
Step 9	<code>switchport mode trunk</code>	Configure the interface as a VLAN trunk port.
Step 10	<code>dot1x pae supplicant</code>	Configure the interface as a port access entity (PAE) supplicant.
Step 11	<code>dot1x credentials <i>profile-name</i></code>	Attach the 802.1x credentials profile to the interface.
Step 12	<code>end</code>	Return to privileged EXEC mode.
Step 13	<code>show running-config interface <i>interface-id</i></code>	Verify your configuration.
Step 14	<code>copy running-config startup-config</code>	(Optional) Save your entries in the configuration file.

This example shows how to configure a switch as a supplicant:

```
Switch# configure terminal
Switch(config)# cisp enable
Switch(config)# dot1x credentials test
Switch(config)# username suppswitch
Switch(config)# password myswitch
Switch(config)# dot1x supplicant force-multicast
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport mode trunk
Switch(config-if)# dot1x pae supplicant
Switch(config-if)# dot1x credentials test
Switch(config-if)# end
```

Configuring NEAT with Auto Smartports Macros

You can also use an Auto Smartports user-defined macro instead of the switch VSA to configure the authenticator switch. For information, see the *Auto Smartports Configuration Guide*.

Configuring 802.1x Authentication with Downloadable ACLs and Redirect URLs

In addition to configuring 802.1x authentication on the switch, you need to configure the ACS. For more information, see the [Cisco Secure ACS configuration guides](#).



Note

You must configure a downloadable ACL on the ACS before downloading it to the switch.

After authentication on the port, you can use the **show ip access-list** privileged EXEC command to display the downloaded ACLs on the port.

Configuring Downloadable ACLs

The policies take effect after client authentication and the client IP address addition to the IP device tracking table. The switch then applies the downloadable ACL to the port.

Beginning in privileged EXEC mode:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip device tracking	Configure the ip device tracking table.
Step 3	aaa new-model	Enables AAA.
Step 4	aaa authorization network default group radius	Sets the authorization method to local. To remove the authorization method, use the no aaa authorization network default group radius command.
Step 5	radius-server vsa send authentication	Configure the radius vsa send authentication.
Step 6	interface <i>interface-id</i>	Specify the port to be configured, and enter interface configuration mode.
Step 7	ip access-group <i>acl-id</i> in	Configure the default ACL on the port in the input direction. Note The <i>acl-id</i> is an access list name or number.
Step 8	show running-config interface <i>interface-id</i>	Verify your configuration.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Configuring a Downloadable Policy

Beginning in privileged EXEC mode:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	access-list <i>access-list-number</i> deny source <i>source-wildcard</i> log	<p>Defines the default port ACL by using a source address and wildcard.</p> <p>The <i>access-list-number</i> is a decimal number from 1 to 99 or 1300 to 1999.</p> <p>Enter deny or permit to specify whether to deny or permit access if conditions are matched.</p> <p>The <i>source</i> is the source address of the network or host that sends a packet, such as this:</p> <ul style="list-style-type: none"> The 32-bit quantity in dotted-decimal format. The keyword any as an abbreviation for source and source-wildcard value of 0.0.0.0 255.255.255.255. You do not need to enter a source-wildcard value. The keyword host as an abbreviation for source and source-wildcard of source 0.0.0.0. <p>(Optional) Applies the source-wildcard wildcard bits to the source.</p> <p>(Optional) Enters log to cause an informational logging message about the packet that matches the entry to be sent to the console.</p>
Step 3	interface <i>interface-id</i>	Enter interface configuration mode.
Step 4	ip access-group <i>acl-id</i> in	<p>Configure the default ACL on the port in the input direction.</p> <p>Note The <i>acl-id</i> is an access list name or number.</p>
Step 5	exit	Returns to global configuration mode.
Step 6	aaa new-model	Enables AAA.
Step 7	aaa authorization network default group radius	Sets the authorization method to local. To remove the authorization method, use the no aaa authorization network default group radius command.
Step 8	ip device tracking	<p>Enables the IP device tracking table.</p> <p>To disable the IP device tracking table, use the no ip device tracking global configuration commands.</p>
Step 9	ip device tracking probe [count interval use-svi]	<p>(Optional) Configures the IP device tracking table:</p> <ul style="list-style-type: none"> count <i>count</i>—Sets the number of times that the switch sends the ARP probe. The range is from 1 to 5. The default is 3. interval <i>interval</i>—Sets the number of seconds that the switch waits for a response before resending the ARP probe. The range is from 30 to 300 seconds. The default is 30 seconds. use-svi—Uses the switch virtual interface (SVI) IP address as source of ARP probes.
Step 10	radius-server vsa send authentication	<p>Configures the network access server to recognize and use vendor-specific attributes.</p> <p>Note The downloadable ACL must be operational.</p>
Step 11	end	Returns to privileged EXEC mode.

	Command	Purpose
Step 12	show ip device tracking all	Displays information about the entries in the IP device tracking table.
Step 13	copy running-config startup-config	(Optional) Saves your entries in the configuration file.

This example shows how to configure a switch for a downloadable policy:

```
Switch# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# aaa new-model
Switch(config)# aaa authorization network default group radius
Switch(config)# ip device tracking
Switch(config)# ip access-list extended default_acl
Switch(config-ext-nacl)# permit ip any any
Switch(config-ext-nacl)# exit
Switch(config)# radius-server vsa send authentication
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# ip access-group default_acl in
Switch(config-if)# exit
```

Configuring VLAN ID-based MAC Authentication

Beginning in privileged EXEC mode, follow these steps:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mab request format attribute 32 vlan access-vlan	Enable VLAN ID-based MAC authentication.
Step 3	copy running-config startup-config	(Optional) Save your entries in the configuration file.

There is no show command to confirm the status of VLAN ID-based MAC authentication. You can use the **debug radius accounting** privileged EXEC command to confirm the RADIUS attribute 32. For more information about this command, see the *Cisco IOS Debug Command Reference*:

http://www.cisco.com/en/US/docs/ios/debug/command/reference/db_q1.html#wp1123741

This example shows how to globally enable VLAN ID-based MAC authentication on a switch:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# mab request format attribute 32 vlan access-vlan
Switch(config-if)# exit
```

Configuring Flexible Authentication Ordering

Beginning in privileged EXEC mode, follow these steps:

	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.

	Command	Purpose
Step 3	authentication order [dot1x mab] {webauth}	(Optional) Set the order of authentication methods used on a port.
Step 4	authentication priority [dot1x mab] {webauth}	(Optional) Add an authentication method to the port-priority list.
Step 5	show authentication	(Optional) 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 attempt 802.1x authentication first, followed by web authentication as fallback method:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/0/2
Switch(config)# authentication order dot1x webauth
```

Configuring Open1x

Beginning in privileged EXEC mode:

	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	authentication control-direction {both in}	(Optional) Configure the port control as unidirectional or bidirectional.
Step 4	authentication fallback <i>name</i>	(Optional) Configure a port to use web authentication as a fallback method for clients that do not support 802.1x authentication.
Step 5	authentication host-mode [multi-auth multi-domain multi-host single-host]	(Optional) Set the authorization manager mode on a port.
Step 6	authentication open	(Optional) Enable or disable open access on a port.
Step 7	authentication order [dot1x mab] {webauth}	(Optional) Set the order of authentication methods used on a port.
Step 8	authentication periodic	(Optional) Enable or disable reauthentication on a port.
Step 9	authentication port-control {auto force-authorized force-un authorized}	(Optional) Enable manual control of the port authorization state.
Step 10	show authentication	(Optional) Verify your entries.
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to configure open 1x on a port:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet1/0/2
Switch(config)# authentication control-direction both
Switch(config)# authentication fallback profile1
Switch(config)# authentication host-mode multi-auth
Switch(config)# authentication open
```



```
Switch(config)# authentication order dot1x webauth
Switch(config)# authentication periodic
Switch(config)# authentication port-control auto
```

Disabling 802.1x Authentication on the Port

You can disable 802.1x authentication on the port by using the **no dot1x pae** interface configuration command.

Beginning in privileged EXEC mode, follow these steps to disable 802.1x authentication on the 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 to be configured, and enter interface configuration mode.
Step 3	no dot1x pae	Disable 802.1x authentication on the port.
Step 4	end	Return to privileged EXEC mode.
Step 5	show authentication interface <i>interface-id</i>	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To configure the port as an 802.1x port access entity (PAE) authenticator, which enables 802.1x on the port but does not allow clients connected to the port to be authorized, use the **dot1x pae authenticator** interface configuration command.

This example shows how to disable 802.1x authentication on the port:

```
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# no dot1x pae authenticator
```

Resetting the 802.1x Authentication Configuration to the Default Values

Beginning in privileged EXEC mode, follow these steps to reset the 802.1x authentication 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>	Enter interface configuration mode, and specify the port to be configured.
Step 3	dot1x default	Reset the 802.1x parameters to the default values.
Step 4	end	Return to privileged EXEC mode.
Step 5	show authentication 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 [details | statistics | summary]** 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.

Beginning with Cisco IOS Release 12.2(55)SE, you can use the **no dot1x logging verbose** global configuration command to filter verbose 802.1x authentication messages. See the [“Authentication Manager CLI Commands”](#) section on page 1-9.

For detailed information about the fields in these displays, see the command reference for this release.



CHAPTER 1

Configuring Interface Characteristics

This chapter defines the types of Catalyst 2960 and 2960-S interfaces and describes how to configure them.

- [Understanding Interface Types, page 1-1](#)
- [Using the Switch USB Ports \(Catalyst 2960-S Switches Only\), page 1-11](#)
- [Using Interface Configuration Mode, page 1-15](#)
- [Using the Ethernet Management Port \(Catalyst 2960-S Only\), page 1-20](#)
- [Configuring Ethernet Interfaces, page 1-23](#)
- [Configuring Layer 3 SVIs, page 1-35](#)
- [Configuring the System MTU, page 1-36](#)
- [Monitoring and Maintaining the Interfaces, page 1-37](#)



Note

For complete syntax and usage information for the commands used in this chapter, see the switch command reference for this release and the *Cisco IOS Interface Command Reference, Release 12.4* on Cisco.com.

Understanding Interface Types

This section describes the different types of supported interfaces with references to chapters that contain more detailed information about configuring these interfaces.



Note

The stack ports on the rear of the switch are not Ethernet ports and cannot be configured.

- [Port-Based VLANs, page 1-2](#)
- [Switch Ports, page 1-2](#)
- [Switch Virtual Interfaces, page 1-3](#)
- [EtherChannel Port Groups, page 1-4](#)
- [Dual-Purpose Uplink Ports, page 1-4](#)
- [Power over Ethernet Ports, page 1-5](#)
- [Connecting Interfaces, page 1-10](#)

Port-Based VLANs

**Note**

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

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 the [Chapter 1, “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 you configure a local port 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 VLANs, use the **vlan** *vlan-id* global configuration command to enter VLAN configuration mode. The VLAN configurations for normal-range VLANs (VLAN IDs 1 to 1005) are saved in the VLAN database. If VTP is version 1 or 2, you must first set VTP mode to transparent to configure extended-range VLANs (VLAN IDs 1006 to 4094). Extended-range VLANs created in transparent mode are not added to the VLAN database but are saved in the switch running configuration. With VTP version 3, you can create extended-range VLANs in client or server mode. These VLANs are saved in the VLAN database.

VLANs can be formed with ports across the stack. The VLAN database is downloaded to all switches in a stack, and all switches in the stack build the same VLAN database. The running configuration and the saved configuration are the same for all switches in a stack.

**Note**

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

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.

Switch Ports

Switch ports are Layer 2-only interfaces associated with a physical port. Switch ports belong to one or more VLANs. You use switch ports for managing the physical interface and associated Layer 2 protocols.

A switch port can be an access port or a trunk port. You can configure a port as an access port or trunk port or let the Dynamic Trunking Protocol (DTP) operate on a per-port basis to set the switchport mode by negotiating with the port on the other end of the link.

Configure switch ports by using the **switchport** interface configuration commands.

**Note**

When you change a Layer 3 interface into Layer 2 mode, the configuration information related to the affected interface might be lost, and the interface is returned to its default configuration.

For detailed information about configuring access port and trunk port characteristics, see [Chapter 1, “Configuring VLANs.”](#)

Access Ports

An access port belongs to and carries the traffic of only one VLAN (unless it is configured as a voice VLAN port). Traffic is received and sent in native formats with no VLAN tagging. Traffic arriving on an access port is assumed to belong to the VLAN assigned to the port.

If an access port receives an 802.1Q tagged packet, the packet is dropped, and the source address is not learned.

Supported access ports:

- Static access ports are manually assigned to a VLAN (or through a RADIUS server for use with IEEE 802.1x). For more information, see the [“802.1x Authentication with VLAN Assignment” section on page 1-17.](#)
- VLAN membership of dynamic access ports is learned through incoming packets. By default, a dynamic access port is not a member of any VLAN. Traffic forwarding to and from the port is enabled only when the port VLAN membership 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 6500 series switch. The Catalyst 2960 or 2960-S 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 1, “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.

The switch supports only 802.1Q trunk ports. An 802.1Q trunk port supports simultaneous tagged and untagged traffic. The 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 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 affects only the associated trunk port. By default, all possible VLANs (VLAN ID 1 to 4094) are in the allowed list. A trunk port can become a member of a VLAN only if VTP knows of the VLAN and if the VLAN is enabled. If VTP learns of a new, enabled VLAN and the VLAN is in the allowed list, the trunk port automatically becomes a member of that VLAN. Traffic is forwarded to and from the trunk port for that VLAN. If VTP learns of an enabled VLAN that is not in the allowed list for a trunk port, the port does not become a member of the VLAN, and traffic for the VLAN is not forwarded to or from the port.

For more information about trunk ports, see [Chapter 1, “Configuring VLANs.”](#)

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. You can associate only one SVI with a VLAN. You configure an SVI for a VLAN only to route 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.



Note You cannot delete interface VLAN 1.

SVIs provide IP host connectivity only to the system. Beginning with Cisco IOS release 12.2(55)SE, you can enable routing and configure static routes on SVIs.



Note Static routing is supported on SVIs only when the switch is running the LAN base image.

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 encapsulated trunk port 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 1-14.



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

EtherChannel Port Groups

EtherChannel port groups treat multiple switch ports as one switch port. An EtherChannel port group acts as a single logical port for high-bandwidth connections between switches or between switches and servers. An EtherChannel balances the traffic load across the links in the channel. If a link within the EtherChannel fails, traffic previously carried over the failed link changes to the remaining links. You can group multiple trunk ports into one logical trunk port or multiple access ports into one logical access port. Most protocols operate over either single ports or aggregated switch ports and do not recognize the physical ports within the port group. The DTP, the Cisco Discovery Protocol (CDP), and the Port Aggregation Protocol (PAgP) operate only on physical ports.

When you configure an EtherChannel, you create a port-channel logical interface and assign an interface to the EtherChannel. 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 1, “Configuring EtherChannels and Link-State Tracking.”](#)

Dual-Purpose Uplink Ports



Note Catalyst 2960-S switches do not have dual-purpose uplink ports.

Some switches support dual-purpose uplink ports. Each uplink port is considered as a single interface with dual front ends—an RJ-45 connector and a small form-factor pluggable (SFP) module connector. The dual front ends are not redundant interfaces, and the switch activates only one connector of the pair.

By default, the switch dynamically selects the interface type that first links up. However, you can use the **media-type** interface configuration command to manually select the RJ-45 connector or the SFP module connector. For information about configuring speed and duplex settings for a dual-purpose uplink, see the [“Setting the Interface Speed and Duplex Parameters”](#) section on page 1-27.

Each uplink port has two LEDs: one shows the status of the RJ-45 port, and one shows the status of the SFP module port. The port LED is on for whichever connector is active. For more information about the LEDs, see the hardware installation guide.

Power over Ethernet Ports



Note

PoE is supported only when the switch is running the LAN base image. Power over Ethernet Plus (PoE+) is supported only on Catalyst 2960-S switches.

PoE switch ports automatically supply power to these connected devices (if the switch senses that there is no power on the circuit):

- Cisco pre-standard powered devices (such as Cisco IP Phones and Cisco Aironet access points)
- IEEE 802.3 af-compliant powered devices
- IEEE 802.3 at-compliant powered devices (PoE+ on Catalyst 2960-S switches only)

A powered device can receive redundant power when it is connected only to a PoE switch port and to an AC power source. After the switch detects a powered device, it determines the device power requirements and then grants or denies power to the device. The switch can also sense the real-time power consumption of the device by monitoring and policing the power usage.

This section has this PoE information:

- [Supported Protocols and Standards, page 1-5](#)
- [Powered-Device Detection and Initial Power Allocation, page 1-6](#)
- [Power Management Modes, page 1-7](#)
- [Power Monitoring and Power Policing, page 1-8](#)

Supported Protocols and Standards

The switch uses these protocols and standards to support PoE:

- CDP with power consumption—The powered device notifies the switch of the amount of power it is consuming. The switch does not reply to the power-consumption messages. The switch can only supply power to or remove power from the PoE port.
- Cisco intelligent power management—The powered device and the switch negotiate through power-negotiation CDP messages for an agreed power-consumption level. The negotiation allows a high-power Cisco powered device, which consumes more than 7 W, to operate at its highest power mode. The powered device first boots up in low-power mode, consumes less than 7 W, and negotiates to obtain enough power to operate in high-power mode. The device changes to high-power mode only when it receives confirmation from the switch.

High-power devices can operate in low-power mode on switches that do not support power-negotiation CDP.

Cisco intelligent power management is backward-compatible with CDP with power consumption; the switch responds according to the CDP message that it receives. CDP is not supported on third-party powered devices; therefore, the switch uses the IEEE classification to determine the power usage of the device.

- IEEE 802.3af—The major features of this standard are powered-device discovery, power administration, disconnect detection, and optional powered-device power classification. For more information, see the standard.
- IEEE 802.3at (Catalyst 2960-S only)—This PoE+ standard supports all the features of 802.1af and increases the maximum power available on each PoE port from 15.4 W to 30 W.

Powered-Device Detection and Initial Power Allocation

The switch detects a Cisco pre-standard or an IEEE-compliant powered device when the PoE-capable port is in the no-shutdown state, PoE is enabled (the default), and the connected device is not being powered by an AC adaptor.

After device detection, the switch determines the device power requirements based on its type:

- A Cisco pre-standard powered device does not provide its power requirement when the switch detects it, so a Catalyst 2960 switch allocates 15.4 W as the initial allocation for power budgeting; a Catalyst 2960-S switch allocates 30 W (PoE+).

The initial power allocation is the maximum amount of power that a powered device requires. The switch initially allocates this amount of power when it detects and powers the powered device. As the switch receives CDP messages from the powered device and as the powered device negotiates power levels with the switch through CDP power-negotiation messages, the initial power allocation might be adjusted.

- The switch classifies the detected IEEE device within a power consumption class. Based on the available power in the power budget, the switch determines if a port can be powered. [Table 1-1](#) lists these levels.

Table 1-1 IEEE Power Classifications

Class	Maximum Power Level Required from the Switch
0 (class status unknown)	15.4 W
1	4 W
2	7 W
3	15.4 W
4	30 W (Catalyst 2960-S only)

The switch monitors and tracks requests for power and grants power only when it is available. The switch tracks its power budget (the amount of power available on the switch for PoE). The switch performs power-accounting calculations when a port is granted or denied power to keep the power budget up to date.

After power is applied to the port, the switch uses CDP to determine the *actual* power consumption requirement of the connected Cisco powered devices, and the switch adjusts the power budget accordingly. This does not apply to third-party PoE devices. The switch processes a request and either grants or denies power. If the request is granted, the switch updates the power budget. If the request is denied, the switch ensures that power to the port is turned off, generates a syslog message, and updates the LEDs. Powered devices can also negotiate with the switch for more power.

If the switch detects a fault caused by an undervoltage, overvoltage, overtemperature, oscillator-fault, or short-circuit condition, it turns off power to the port, generates a syslog message, and updates the power budget and LEDs.

In a Catalyst 2960-S switch stack, the PoE feature operates the same whether or not the switch is a stack member. The power budget is per-switch and independent of any other switch in the stack. Election of a new stack master does not affect PoE operation. The stack master keeps track of PoE status for all switches and ports in the stack and includes the status in output displays.

Power Management Modes

Supported PoE modes:

- **auto**—The switch automatically detects if the connected device requires power. If the switch discovers a powered device connected to the port and if the switch has enough power, it grants power, updates the power budget, turns on power to the port on a first-come, first-served basis, and updates the LEDs. For LED information, see the hardware installation guide.

If the switch has enough power for all the powered devices, they all come up. If enough power is available for all powered devices connected to the switch, power is turned on to all devices. If there is not enough available PoE, or if a device is disconnected and reconnected while other devices are waiting for power, it cannot be determined which devices are granted or are denied power.

If granting power would exceed the system power budget, the switch denies power, ensures that power to the port is turned off, generates a syslog message, and updates the LEDs. After power has been denied, the switch periodically rechecks the power budget and continues to attempt to grant the request for power.

If a device being powered by the switch is then connected to wall power, the switch might continue to power the device. The switch might continue to report that it is still powering the device whether the device is being powered by the switch or receiving power from an AC power source.

If a powered device is removed, the switch automatically detects the disconnect and removes power from the port. You can connect a nonpowered device without damaging it.

You can specify the maximum wattage that is allowed on the port. If the IEEE class maximum wattage of the powered device is greater than the configured maximum value, the switch does not provide power to the port. If the switch powers a powered device, but the powered device later requests through CDP messages more than the configured maximum value, the switch removes power to the port. The power that was allocated to the powered device is reclaimed into the global power budget. If you do not specify a wattage, the switch delivers the maximum value. Use the **auto** setting on any PoE port. The auto mode is the default setting.

- **static**—The switch pre-allocates power to the port (even when no powered device is connected) and guarantees that power will be available for the port. The switch allocates the port configured maximum wattage, and the amount is never adjusted through the IEEE class or by CDP messages from the powered device. Because power is pre-allocated, any powered device that uses less than or equal to the maximum wattage is guaranteed to be powered when it is connected to the static port. The port no longer participates in the first-come, first-served model.

However, if the powered-device IEEE class is greater than the maximum wattage, the switch does not supply power to it. If the switch learns through CDP messages that the powered device needs more than the maximum wattage, the powered device is shutdown.

If you do not specify a wattage, the switch pre-allocates the maximum value. The switch powers the port only if it discovers a powered device. Use the **static** setting on a high-priority interface.

- **never**—The switch disables powered-device detection and never powers the PoE port even if an unpowered device is connected. Use this mode only when you want to make sure power is never applied to a PoE-capable port, making the port a data-only port.

For information on configuring a PoE port, see the [“Configuring a Power Management Mode on a PoE Port” section on page 1-30](#).

Power Monitoring and Power Policing

When policing of the real-time power consumption is enabled, the switch takes action when a powered device consumes more power than the maximum amount allocated, also referred to as the *cutoff-power value*.

When PoE is enabled, the switch senses the real-time power consumption of the powered device and monitors the power consumption of the connected powered device; this is called *power monitoring* or *power sensing*. The switch also uses the *power policing* feature to police the power usage.

Power monitoring is backward-compatible with Cisco intelligent power management and CDP-based power consumption. It works with these features to ensure that the PoE port can supply power to the powered device. For more information about these PoE features, see the [“Powered-Device Detection and Initial Power Allocation” section on page 1-6](#).

The switch senses the power consumption of the connected device as follows:

1. The switch monitors the real-time power consumption on individual ports.
2. The switch records the power consumption, including peak power usage, and reports the information through an SNMP MIB, CISCO-POWER-ETHERNET-EXT-MIB.
3. If power policing is enabled, the switch polices power usage by comparing the real-time power consumption to the maximum power allocated to the device. For more information about the maximum power consumption, also referred to as the *cutoff power*, on a PoE port, see the [“Maximum Power Allocation \(Cutoff Power\) on a PoE Port” section on page 1-8](#).

If the device uses more than the maximum power allocation on the port, the switch can either turn off power to the port, or the switch can generate a syslog message and update the LEDs (the port LED is now blinking amber) while still providing power to the device based on the switch configuration. By default, power-usage policing is disabled on all PoE ports.

If error recovery from the PoE error-disabled state is enabled, the switch automatically takes the PoE port out of the error-disabled state after the specified amount of time.

If error recovery is disabled, you can manually re-enable the PoE port by using the **shutdown** and **no shutdown** interface configuration commands.

4. If policing is disabled, no action occurs when the powered device consumes more than the maximum power allocation on the PoE port, which could adversely affect the switch.

Maximum Power Allocation (Cutoff Power) on a PoE Port

When power policing is enabled, the switch determines the cutoff power on the PoE port in this order:

1. Manually when you set the user-defined power level that the switch budgets for the port by using the **power inline consumption default** *wattage* global or interface configuration command
2. Manually when you set the user-defined power level that limits the power allowed on the port by using the **power inline auto max** *max-wattage* or the **power inline static max** *max-wattage* interface configuration command
3. Automatically when the switch sets the power usage of the device by using CDP power negotiation or by the IEEE classification and LLDP power negotiation.

Use the first or second method in the previous list to manually configure the cutoff-power value by entering the **power inline consumption default** *wattage* or the **power inline [auto | static max]** *max-wattage* command. If you do not manually configure the cutoff-power value, the switch automatically determines the value by using CDP power negotiation. If the switch cannot determine the value by using one of these methods, it uses the default value of 15.4 W. On a Catalyst 2960-S switch with PoE+, if you do not manually configure the cutoff-power value, the switch automatically

determines it by using CDP power negotiation or the device IEEE classification and LLDP power negotiation. If CDP or LLDP are not enabled, the default value of 30 W is applied. However without CDP or LLDP, the switch does not allow devices to consume more than 15.4 W of power because values from 15400 to 30000 mW are only allocated based on CDP or LLDP requests. If a powered device consumes more than 15.4 W without CDP or LLDP negotiation, the device might be in violation of the maximum current (I_{max}) limitation and might experience an I_{cut} fault for drawing more current than the maximum. The port remains in the fault state for a time before attempting to power on again. If the port continuously draws more than 15.4 W, the cycle repeats.

**Note**

When a powered device connected to a PoE+ port restarts and sends a CDP or LLDP packet with a power TLV, the switch locks to the power-negotiation protocol of that first packet and does not respond to power requests from the other protocol. For example, if the switch is locked to CDP, it does not provide power to devices that send LLDP requests. If CDP is disabled after the switch has locked on it, the switch does not respond to LLDP power requests and can no longer power on any accessories. In this case, you should restart the powered device.

Power Consumption Values

You can configure the initial power allocation and the maximum power allocation on a port. However, these values are only the configured values that determine when the switch should turn on or turn off power on the PoE port. The maximum power allocation is not the same as the actual power consumption of the powered device. The actual cutoff power value that the switch uses for power policing is not equal to the configured power value.

When power policing is enabled, the switch polices the power usage *at the switch port*, which is greater than the power consumption of the device. When you manually set the maximum power allocation, you must consider the power loss over the cable from the switch port to the powered device. The cutoff power is the sum of the rated power consumption of the powered device and the worst-case power loss over the cable.

The actual amount of power consumed by a powered device on a PoE port is the cutoff-power value plus a calibration factor of 500 mW (0.5 W). The actual cutoff value is approximate and varies from the configured value by a percentage of the configured value. For example, if the configured cutoff power is 12 W, the actual cutoff-value is 11.4 W, which is 0.05% less than the configured value.

We recommend that you enable power policing when PoE is enabled on your switch. For example, if policing is disabled and you set the cutoff-power value by using the **power inline auto max 6300** interface configuration command, the configured maximum power allocation on the PoE port is 6.3 W (6300 mW). The switch provides power to the connected devices on the port if the device needs up to 6.3 W. If the CDP-power negotiated value or the IEEE classification value exceeds the configured cutoff value, the switch does not provide power to the connected device. After the switch turns on power to the PoE port, the switch does not police the real-time power consumption of the device, and the device can consume more power than the maximum allocated amount, which could adversely affect the switch and the devices connected to the other PoE ports.

Because the switch supports internal power supplies and the Cisco Redundant Power System 2300 (also referred to as the RPS 2300), the total amount of power available for the powered devices varies depending on the power supply configuration.

- If a power supply is removed and replaced by a new power supply with less power and the switch does not have enough power for the powered devices, the switch denies power to the PoE ports that are in auto mode in descending order of the port numbers. If the switch still does not have enough power, it denies power to the PoE ports in static mode in descending order of the port numbers.

- If the new power supply supports more power than the previous one and the switch now has more power available, the switch grants power to the PoE ports in static mode in ascending order of the port numbers. If it still has power available, the switch then grants power to the PoE ports in auto mode in ascending order of the port numbers.


For configuration information, see the [“Configuring Power Policing”](#) section on page 1-33.

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.


In the configuration shown in [Figure 1-1](#), when Host A in VLAN 20 sends data to Host B in VLAN 30, the data must go from Host A to the switch, to the router, back to the switch, and then to Host B.

Figure 1-1 Connecting VLANs with Layer 2 Switches



With a standard Layer 2 switch, ports in different VLANs have to exchange information through a router. By using the switch with routing enabled, when you configure both VLAN 20 and VLAN 30 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 1-2](#)).

Figure 1-2 Connecting VLANs with a Layer 3 Switch



Using the Switch USB Ports (Catalyst 2960-S Switches Only)

The Catalyst 2960-S switch has two USB ports on the front panel:

- [USB Mini-Type B Console Port, page 1-11](#)
- [USB Type A Port, page 1-14](#)

USB Mini-Type B Console Port

The switch has two console ports available—a USB mini-Type B console connection and an RJ-45 console port. Console *output* appears on devices connected to both ports, but console *input* is active on only one port at a time. The USB connector takes precedence over the RJ-45 connector.



Note

Windows PCs require a driver for the USB port. See the hardware installation guide for driver installation instructions.

Use the supplied USB Type A-to-USB mini-Type B cable to connect a PC or other device to the switch. The connected device must include a terminal emulation application. When the switch detects a valid USB connection to a powered-on device that supports host functionality (such as a PC), input from the RJ-45 console is immediately disabled, and input from the USB console is enabled. Removing the USB connection immediately reenables input from the RJ-45 console connection. An LED on the switch shows which console connection is in use.

Console Port Change Logs

At software startup, a log shows whether the USB or the RJ-45 console is active. Each switch in a stack issues this log. Every switch always first displays the RJ-45 media type.

In the sample output, switch 1 has a connected USB console cable. Because the bootloader did not change to the USB console, the first log from switch 1 shows the RJ-45 console. A short time later, the console changes and the USB console log appears. Switch 2 and switch 3 have connected RJ-45 console cables.

```
switch-stack-1
*Mar 1 00:01:00.171: %USB_CONSOLE-6-MEDIA_RJ45: Console media-type is RJ45.
*Mar 1 00:01:00.431: %USB_CONSOLE-6-MEDIA_USB: Console media-type is USB.

switch-stack-2
*Mar 1 00:01:09.835: %USB_CONSOLE-6-MEDIA_RJ45: Console media-type is RJ45.
switch-stack-3)
*Mar 1 00:01:10.523: %USB_CONSOLE-6-MEDIA_RJ45: Console media-type is RJ45.
```

When the USB cable is removed or the PC de-activates the USB connection, the hardware automatically changes to the RJ-45 console interface:

```
switch-stack-1
Mar 1 00:20:48.635: %USB_CONSOLE-6-MEDIA_RJ45: Console media-type is RJ45.
```

You can configure the console type to always be RJ-45, and you can configure an inactivity timeout for the USB connector.

Configuring the Console Media Type

Beginning in privileged EXEC mode, follow these steps to select the RJ-45 console media type. If you configure the RJ-45 console, USB console operation is disabled, and input always remains with the RJ-45 console.

This configuration applies to all switches in a stack.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	line console 0	Configure the console. Enter line configuration mode.
Step 3	media-type rj45	Configure the console media type to always be RJ-45. If you do not enter this command and both types are connected, the default is USB.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-configuration	Verify your settings.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example disables the USB console media type and enables the RJ-45 console media type.

```
Switch# configure terminal
Switch(config)# line console 0
Switch(config-line)# media-type rj45
```

This configuration terminates any active USB console media type in the stack. A log shows that this termination has occurred. This example shows that the console on switch 1 reverted to RJ-45.

```
*Mar 1 00:25:36.860: %USB_CONSOLE-6-CONFIG_DISABLE: Console media-type USB disabled by
system configuration, media-type reverted to RJ45.
```

At this point no switches in the stack allow a USB console to have input. A log entry shows when a console cable is attached. If a USB console cable is connected to switch 2, it is prevented from providing input.

```
*Mar 1 00:34:27.498: %USB_CONSOLE-6-CONFIG_DISALLOW: Console media-type USB is disallowed
by system configuration, media-type remains RJ45. (switch-stk-2)
```

This example reverses the previous configuration and immediately activates any USB console that is connected.

```
Switch# configure terminal
Switch(config)# line console 0
Switch(config-line)# no media-type rj45
```

Configuring the USB Inactivity Timeout

The configurable inactivity timeout reactivates the RJ-45 console port if the USB console port is activated but no input activity occurs on it for a specified time period. When the USB console port is deactivated due to a timeout, you can restore its operation by disconnecting and reconnecting the USB cable.



Note

The configured inactivity timeout applies to all switches in a stack. However, a timeout on one switch does *not* cause a timeout on other switches in the stack.

Beginning in privileged EXEC mode, follow these steps to configure an inactivity timeout.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	line console 0	Configure the console port. Enter console line configuration mode.
Step 3	usb-inactivity-timeout <i>timeout-minutes</i>	Specify an inactivity timeout for the console port. The range is 1 to 240 minutes. The default is to have no timeout configured.
Step 4	show running-configuration	Verify your setting.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example configures the inactivity timeout to 30 minutes:

```
Switch# configure terminal
Switch#(config)# line console 0
Switch#(config-line)# usb-inactivity-timeout 30
```

To disable the configuration, use these commands:

```
Switch#(config)# line console 0
Switch#(config-line)# no usb-inactivity-timeout
```

If there is no (input) activity on a USB console port for the configured number of minutes, the inactivity timeout setting applies to the RJ-45 port, and a log shows this occurrence:

```
*Mar 1 00:47:25.625: %USB_CONSOLE-6-INACTIVITY_DISABLE: Console media-type USB disabled
due to inactivity, media-type reverted to RJ45.
```

At this point, the only way to reactivate the USB console port is to disconnect and reconnect the cable.

When the USB cable on the switch has been disconnected and reconnected, a log similar to this appears:

```
*Mar 1 00:48:28.640: %USB_CONSOLE-6-MEDIA_USB: Console media-type is USB.
```

USB Type A Port

The USB Type A port provides access to external USB flash devices, also known as thumb drives or USB keys. The switch supports Cisco 64 MB, 256 MB, 512 MB and 1 GB flash drives. You can use standard Cisco IOS command-line interface (CLI) commands to read, write, erase, and copy to or from the flash device. You can also configure the switch to boot from the USB flash drive.

Beginning in privileged EXEC mode, follow these steps to allow booting from the USB flash device.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	boot system flash usbflash0: <i>image</i>	Configure the switch to boot from the USB flash device. The <i>image</i> is the name of the bootable image.
Step 3	show running-configuration	Verify your setting.
Step 4	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To get information about the USB device, use the **show usb {controllers | device | driver | port | tree}** privileged EXEC command.

This example configures the switch to boot from the Catalyst 2960-S flash device. The image is the Catalyst 2960-S LAN base image.

```
Switch# configure terminal
Switch#(config)# boot system flash usbflash0: c2960s-lanbase-mz
```

To disable booting from flash, enter the **no** form of the command.

This is sample output from the **show usb device** command:

```
Switch# show usb device
Host Controller: 1
Address: 0x1
Device Configured: YES
Device Supported: YES
Description: STEC USB 1GB
Manufacturer: STEC
Version: 1.0
Serial Number: STI 3D508232204731
Device Handle: 0x1010000
USB Version Compliance: 2.0
Class Code: 0x0
Subclass Code: 0x0
Protocol: 0x0
Vendor ID: 0x136b
Product ID: 0x918
Max. Packet Size of Endpoint Zero: 64
Number of Configurations: 1
Speed: High
Selected Configuration: 1
Selected Interface: 0
```



```
Configuration:
  Number: 1
  Number of Interfaces: 1
  Description: Storage
  Attributes: None
  Max Power: 200 mA

Interface:
  Number: 0
  Description: Bulk
  Class Code: 8
  Subclass: 6
  Protocol: 80
  Number of Endpoints: 2

Endpoint:
  Number: 1
  Transfer Type: BULK
  Transfer Direction: Device to Host
  Max Packet: 512
  Interval: 0

Endpoint:
  Number: 2
  Transfer Type: BULK
  Transfer Direction: Host to Device
  Max Packet: 512
  Interval: 0
```

This is sample output from the **show usb port** command:

```
Switch# show usb port
Port Number: 0
Status: Enabled
Connection State: Connected
Speed: High
Power State: ON
```

Using Interface Configuration Mode

The switch supports these interface types:

- Physical ports—switch ports
- VLANs—switch virtual interfaces
- Port channels—EtherChannel interfaces

You can also configure a range of interfaces (see the [“Configuring a Range of Interfaces”](#) section on page 1-17).

To configure a physical interface (port) on a Catalyst 2960 switch or a Catalyst 2960-S switch running the LAN Lite image, specify the interface type, module number, and switch port number, and enter interface configuration mode. To configure a port on a Catalyst 2960-S switch running the LAN base image (supporting stacking), specify the interface type, stack member number, module number, and switch port number, and enter interface configuration mode.

- Type—Port types depend on those supported on the switch. Possible types are: Fast Ethernet (fastethernet or fa) for 10/100 Mb/s Ethernet, Gigabit Ethernet (gigabitethernet or gi) for 10/100/1000 Mb/s Ethernet ports, 10-Gigabit Ethernet (tengigabitethernet or te) for 10,000 Mb/s, or small form-factor pluggable (SFP) module Gigabit Ethernet interfaces.

- Stack member number—The number that identifies the switch within the stack. The switch number range is 1 to 4 and is assigned the first time the switch initializes. The default switch number, before it is integrated into a switch stack, is 1. When a switch has been assigned a stack member number, it keeps that number until another is assigned to it.
- Module number—The module or slot number on the switch (always 0).
- Port number—The interface number on the switch. The port numbers always begin at 1, starting with the far left port when facing the front of the switch, for example, `gigabitethernet1/0/1`. For a switch with 10/100/1000 ports and SFP module ports, SFP module ports are numbered consecutively following the 10/100/1000 ports.

You can identify physical interfaces by looking at the switch. You can also use the **show** privileged EXEC commands to display information about a specific interface or all the interfaces. The remainder of this chapter primarily provides physical interface configuration procedures.

These examples identify interfaces on a Catalyst 2960-S switch running the LAN base image:

- To configure 10/100/1000 port 4 on a standalone switch, enter this command:

```
Switch(config)# interface gigabit
tetherenet1/0/4
```

- To configure 10/100 port 4 on stack member 3, enter this command:

```
Switch(config)# interface gigabitethernet3/0/4
```

This example identifies an interface on a Catalyst 2960 switch or a Catalyst 2960-S switch running the LAN Lite image:

- To configure 10/100/1000 port 4, enter this command:

```
Switch(config)# interface gigabitethernet0/4
```



Note

Configuration examples and outputs in this book might not be specific to your switch, particularly regarding the presence of a stack member number.

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 the interface number, Gigabit Ethernet port 1 in this example:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)#
```



Note

Entering a space between the interface type and interface number is optional

Step 3 Follow each **interface** command with the 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 1-37.

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> }	Specify the range of interfaces (VLANs or physical ports) to be configured, and enter interface-range configuration mode. <ul style="list-style-type: none"> You can use the interface range command to configure up to five port ranges or a previously defined macro. The macro variable is explained in the “Configuring and Using Interface Range Macros” section on page 1-19. 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		Use the normal configuration commands to apply the configuration parameters to all interfaces in the range. Each command is executed as it is entered.
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*, depending on port types on the switch:
 - **vlan** *vlan-ID*, where the VLAN ID is 1 to 4094



Note Although the command-line interface shows options to set multiple VLANs, these options are not supported on Catalyst 2960 and 2960-S switches.

- , where the module is always 0
- **fastethernet** module/{*first port*} - {*last port*}, where the module is always 0
- **gigabitethernet** module/{*first port*} - {*last port*}, where the module is always 0
- **port-channel** *port-channel-number* - *port-channel-number*, where the *port-channel-number* is 1 to 6



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, **interface range gigabitethernet 0/1 - 4** is a valid range; **interface range gigabitethernet0/1-4** is not.

- 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 in a range must be the same type (all Fast Ethernet ports, all Gigabit Ethernet 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 ports 1 to 2 to 100 Mb/s:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet0/1 - 2
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 1 to 3 and Gigabit Ethernet ports 1 and 2 to receive flow-control pause frames:

```
Switch# configure terminal
Switch(config)# interface range fastethernet0/1 - 3, gigabitethernet0/1 - 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 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. <ul style="list-style-type: none"> The <i>macro_name</i> is a 32-character maximum character string. A macro can contain up to five comma-separated interface ranges. 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.
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*, depending on port types on the switch:
 - vlan** *vlan-ID*, where the VLAN ID is 1 to 4094



Note Although the command-line interface shows options to set multiple VLANs, these options are not supported on Catalyst 2960 switches.

- fastethernet** *module*/{*first port*} - {*last port*}, where the module is always 0
- gigabitethernet** *module*/{*first port*} - {*last port*}, where the module is always 0
- port-channel** *port-channel-number* - *port-channel-number*, where the *port-channel-number* is 1 to 6



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 entering an *interface-rang*.

For example, **gigabitethernet 0/1 - 4** is a valid range; **gigabitethernet0/1-4** is not.

- 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 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 named *enet_list* to include ports 1 and 2 and to verify the macro configuration:

```
Switch# configure terminal
Switch(config)# define interface-range enet_list gigabitethernet0/1 - 2
Switch(config)# end
Switch# show running-config | include define
Switch# define interface-range enet_list gigabitethernet0/1 - 2
```

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

```
Switch# configure terminal
Switch(config)# define interface-range macro1 fastethernet0/1 - 2, gigabitethernet0/1 - 2
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#
```

Using the Ethernet Management Port (Catalyst 2960-S Only)



Note

The Ethernet management port is not supported on Catalyst 2960 switches.

- [Understanding the Ethernet Management Port, page 1-21](#)
- [Supported Features on the Ethernet Management Port, page 1-21](#)
- [Configuring the Ethernet Management Port, page 1-22](#)
- [TFTP and the Ethernet Management Port, page 1-22](#)


Understanding the Ethernet Management Port

The Ethernet management port, also referred to as the *Fa0* or *fastethernet0* port, is a Layer 3 host port to which you can connect a PC. You can use the Ethernet management port instead of the switch console port for network management. When managing a switch stack, connect the PC to the Ethernet management port on a Catalyst 2960-S stack member.

When connecting a PC to the Ethernet management port, you must assign an IP address.


For a Catalyst 2960-S standalone switch, connect the Ethernet management port to the PC as shown in [Figure 1-3](#).

Figure 1-3 Connecting a Switch to a PC



In a Catalyst 2960-S stack, all the Ethernet management ports on the stack members are connected to a hub to which the PC is connected. As shown in [Figure 1-4](#), the active link is from the Ethernet management port on the stack master (switch 2) through the hub, to the PC. If the stack master fails and a new stack master is elected, the active link is now from the Ethernet management port on the new stack master to the PC.

Figure 1-4 Connecting a Switch Stack to a PC



By default, the Ethernet management port is enabled.

Supported Features on the Ethernet Management Port

The Ethernet management port supports these features:

- Express Setup (only in switch stacks)
- Network Assistant
- Telnet with passwords

- TFTP
- Secure Shell (SSH)
- DHCP-based autoconfiguration
- SMNP (only the ENTITY-MIB and the IF-MIB)
- IP ping
- Interface features
 - Speed—10 Mb/s, 100 Mb/s, and autonegotiation
 - Duplex mode—Full, half, and autonegotiation
 - Loopback detection
- Cisco Discovery Protocol (CDP)
- DHCP relay agent
- IPv4 and IPv6 access control lists (ACLs)

**Caution**

Before enabling a feature on the Ethernet management port, make sure that the feature is supported. If you try to configure an unsupported feature on the Ethernet Management port, the feature might not work properly, and the switch might fail.

Configuring the Ethernet Management Port

To specify the Ethernet management port in the CLI, enter **fastethernet0**.

To disable the port, use the **shutdown** interface configuration command. To enable the port, use the **no shutdown** interface configuration command.

To find out the link status to the PC, you can monitor the LED for the Ethernet management port. The LED is green (on) when the link is active, and the LED is off when the link is down. The LED is amber when there is a POST failure.

To display the link status, use the **show interfaces fastethernet 0** privileged EXEC command.

TFTP and the Ethernet Management Port

Use the commands in [Table 1-2](#) when using TFTP to download or upload a configuration file to the boot loader.

Table 1-2 Boot Loader Commands

Command	Description
arp [<i>ip_address</i>]	Displays the currently cached ARP ¹ table when this command is entered without the <i>ip_address</i> parameter. Enables ARP to associate a MAC address with the specified IP address when this command is entered with the <i>ip_address</i> parameter.
mgmt_clr	Clears the statistics for the Ethernet management port.
mgmt_init	Starts the Ethernet management port.

Table 1-2 *Boot Loader Commands (continued)*

Command	Description
mgmt_show	Displays the statistics for the Ethernet management port.
ping <i>host_ip_address</i>	Sends ICMP ECHO_REQUEST packets to the specified network host.
boot tftp: <i>/file-url ...</i>	Loads and boots an executable image from the TFTP server and enters the command-line interface. For more details, see the command reference for this release.
copy tftp: <i>/source-file-url filesystem:/destination-file-url</i>	Copies a Cisco IOS image from the TFTP server to the specified location. For more details, see the command reference for this release.

1. ARP = Address Resolution Protocol.

Configuring Ethernet Interfaces

- [Default Ethernet Interface Configuration, page 1-23](#)
- [Setting the Type of a Dual-Purpose Uplink Port, page 1-25](#)
- [Setting the Type of a Dual-Purpose Uplink Port, page 1-25](#)
- [Configuring Interface Speed and Duplex Mode, page 1-26](#)
- [Configuring IEEE 802.3x Flow Control, page 1-28](#)
- [Configuring Auto-MDIX on an Interface, page 1-29](#)
- [Configuring a Power Management Mode on a PoE Port, page 1-30](#)
- [Budgeting Power for Devices Connected to a PoE Port, page 1-31](#)
- [Configuring Power Policing, page 1-33](#)
- [Adding a Description for an Interface, page 1-34](#)

Default Ethernet Interface Configuration

[Table 1-3](#) shows the Ethernet interface default configuration. For more details on the VLAN parameters listed in the table, see [Chapter 1, “Configuring VLANs.”](#) For details on controlling traffic to the port, see [Chapter 1, “Configuring Port-Based Traffic Control.”](#)

Table 1-3 *Default Layer 2 Ethernet Interface Configuration*

Feature	Default Setting
Allowed VLAN range	VLANs 1 to 4094.
Default VLAN (for access ports)	VLAN 1.
Native VLAN (for IEEE 802.1Q trunks)	VLAN 1.
802.1p priority-tagged traffic	Drop all packets tagged with VLAN 0.
VLAN trunking	Switchport mode dynamic auto (supports DTP).

Table 1-3 Default Layer 2 Ethernet Interface Configuration (continued)

Feature	Default Setting
Port enable state	All ports are enabled.
Port description	None defined.
Speed	Autonegotiate.
Duplex mode	Autonegotiate.
Flow control	Flow control is set to receive: off . It is always off for sent packets.
EtherChannel (PAgP)	Disabled on all Ethernet ports. Chapter 1, “Configuring EtherChannels and Link-State Tracking.”
Port blocking (unknown multicast and unknown unicast traffic)	Disabled (not blocked). See the “Configuring Port Blocking” section on page 1-7.
Broadcast, multicast, and unicast storm control	Disabled. See the “Default Storm Control Configuration” section on page 1-3.
Protected port	Disabled. See the “Configuring Protected Ports” section on page 1-6.
Port security	Disabled. See the “Default Port Security Configuration” section on page 1-11.
Port Fast	Disabled. See the “Default Optional Spanning-Tree Configuration” section on page 1-12.
Auto-MDIX	Enabled. Note The switch might not support a pre-standard powered device—such as Cisco IP phones and access points that do not fully support IEEE 802.3af—if that powered device is connected to the switch through a crossover cable. This is regardless of whether auto-MDIX is enabled on the switch port.
Power over Ethernet (PoE)	Enabled (auto).
Keepalive messages	Disabled on SFP module ports; enabled on all other ports.

Setting the Type of a Dual-Purpose Uplink Port



Note Only Catalyst 2960 switches have dual-purpose uplinks ports.

Some switches support dual-purpose uplink ports. By default, the switch dynamically selects the interface type that first links up. However, you can use the **media-type** interface configuration command to manually select the RJ-45 connector or the SFP module connector. For more information, see the “Dual-Purpose Uplink Ports” section on page 1-4.

Beginning in privileged EXEC mode, follow these steps to select which dual-purpose uplink to activate so that you can set the speed and duplex. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the dual-purpose uplink port to be configured, and enter interface configuration mode.
Step 3	media-type { auto-select rj45 sfp }	Select the interface and type of a dual-purpose uplink port. The keywords have these meanings: <ul style="list-style-type: none"> • auto-select—The switch dynamically selects the type. When link up is achieved, the switch disables the other type until the active link goes down. When the active link goes down, the switch enables both types until one of them links up. In auto-select mode, the switch configures both types with autonegotiation of speed and duplex (the default). Depending on the type of installed SFP module, the switch might not be able to dynamically select it. For more information, see the information that follows this procedure. • rj45—The switch disables the SFP module interface. If you connect an SFP module to this port, it cannot attain a link even if the RJ-45 side is down or is not connected. In this mode, the dual-purpose port behaves like a 10/100/1000BASE-TX interface. You can configure the speed and duplex settings consistent with this interface type. • sfp—The switch disables the RJ-45 interface. If you connect a cable to the RJ-45 port, it cannot attain a link even if the SFP module side is down or if the SFP module is not present. Based on the type of installed SFP module, you can configure the speed and duplex settings consistent with this interface type. For information about setting the speed and duplex, see the “Speed and Duplex Configuration Guidelines” section on page 1-26.
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces <i>interface-id</i> transceiver properties	Verify your setting.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **media-type auto interface** or the **no media-type interface** configuration commands.

The switch configures both types to autonegotiate speed and duplex (the default). If you configure **auto-select**, you cannot configure the **speed** and **duplex** interface configuration commands.

When the switch powers on or when you enable a dual-purpose uplink port through the **shutdown** and the **no shutdown** interface configuration commands, the switch gives preference to the SFP module interface. In all other situations, the switch selects the active link based on which type first links up.

The switch operates with 100BASE-*x* (where *x* is -BX, -FX-FE, -LX) SFP modules as follows:

- When the 100BASE-*x* SFP module is inserted into the module slot and there is no link on the RJ-45 side, the switch disables the RJ-45 interface and selects the SFP module interface. This is the behavior even if there is no cable connected and if there is no link on the SFP module side.
- When the 100BASE-*x* SFP module is inserted and there is a link on the RJ-45 side, the switch continues with that link. If the link goes down, the switch disables the RJ-45 side and selects the SFP module interface.
- When the 100BASE-*x* SFP module is removed, the switch again dynamically selects the type (**auto-select**) and re-enables the RJ-45 side.

The switch does not have this behavior with 100BASE-FX-GE SFP modules.

Configuring Interface Speed and Duplex Mode

Depending on the supported port types, Ethernet interfaces on the switch operate at 10, 100, or 1000 Mb/s, or 10,000 Mb/s 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-Mb/s ports operate in half-duplex mode, which means that stations can either receive or send traffic.

Switch models can include combinations of Fast Ethernet (10/100-Mb/s) ports, Gigabit Ethernet (10/100/1000-Mb/s) ports, 10-Gigabit module ports, and small form-factor pluggable (SFP) module slots supporting SFP modules.

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

- [Speed and Duplex Configuration Guidelines, page 1-26](#)
- [Setting the Interface Speed and Duplex Parameters, page 1-27](#)

Speed and Duplex Configuration Guidelines

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

- Fast Ethernet (10/100-Mb/s) ports support all speed and duplex options.
- Gigabit Ethernet (10/100/1000-Mb/s) ports support all speed options and all duplex options (auto, half, and full). However, Gigabit Ethernet ports operating at 1000 Mb/s do not support half-duplex mode.
- For SFP module ports, the speed and duplex CLI options change depending on the SFP module type:
 - The 1000BASE-*x* (where *x* is -BX, -CWDM, -LX, -SX, and -ZX) SFP module ports support the **nonegotiate** keyword in the **speed** interface configuration command. Duplex options are not supported.
 - The 1000BASE-T SFP module ports support the same speed and duplex options as the 10/100/1000-Mb/s ports.

- The 100BASE-*x* (where *x* is -BX, -CWDM, -LX, -SX, and -ZX) SFP module ports support only 100 Mb/s. These modules support full- and half- duplex options but do not support autonegotiation.

For information about which SFP modules are supported on your switch, see the product release notes.

- 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.
- 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>	Specify the physical interface to be configured, and enter interface configuration mode.
Step 3	speed { 10 100 1000 auto [10 100 1000] nonegotiate }	<p>Enter the appropriate speed parameter for the interface:</p> <ul style="list-style-type: none"> • Enter 10, 100, or 1000 to set a specific speed for the interface. The 1000 keyword is available only for 10/100/1000 Mb/s ports. • Enter auto to enable the interface to autonegotiate speed with the connected device. If you use the 10, 100, or the 1000 keywords with the auto keyword, the port autonegotiates only at the specified speeds. • The nonegotiate keyword is available only for SFP module ports. SFP module ports operate only at 1000 Mb/s but can be configured to not negotiate if connected to a device that does not support autonegotiation. <p>For more information about speed settings, see the “Speed and Duplex Configuration Guidelines” section on page 1-26.</p>
Step 4	duplex { auto full half }	<p>Enter the duplex parameter for the interface.</p> <p>Enable half-duplex mode (for interfaces operating only at 10 or 100 Mb/s). You cannot configure half-duplex mode for interfaces operating at 1000 Mb/s.</p> <p>For more information about duplex settings, see the “Speed and Duplex Configuration Guidelines” section on page 1-26.</p>
Step 5	end	Return to privileged EXEC mode.

	Command	Purpose
Step 6	<code>show interfaces interface-id</code>	Display the interface speed and duplex mode configuration.
Step 7	<code>copy running-config startup-config</code>	(Optional) Save your entries in the configuration file.

Use the **no speed** and **no duplex** interface configuration commands to return the interface to the default speed and duplex settings (autonegotiate). To return all interface settings to the defaults, use the **default interface interface-id** interface configuration command.

This example shows how to set the interface speed to 10 Mb/s and the duplex mode to half on a 10/100 Mb/s port:

```
Switch# configure terminal
Switch(config)# interface fasttethernet0/3
Switch(config-if)# speed 10
Switch(config-if)# duplex half
```

This example shows how to set the interface speed to 100 Mb/s on a 10/100/1000 Mb/s port:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# speed 100
```

Configuring IEEE 802.3x 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 by sending a pause frame to stop sending until the condition clears. 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

Ports on the switch can receive, but not send, 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, see 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>	Specify the physical interface to be configured, and enter interface configuration mode.
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 an Interface

When automatic medium-dependent interface crossover (auto-MDIX) is enabled on an interface, the interface 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, see the hardware installation guide.

Auto-MDIX is enabled by default. When you enable auto-MDIX, you must also set the interface speed and duplex to **auto** so that the feature operates correctly.

Auto-MDIX is supported on all 10/100 and 10/100/1000-Mb/s interfaces. It is not supported on 1000BASE-SX or -LX SFP module interfaces.

Table 1-4 shows the link states that result from auto-MDIX settings and correct and incorrect cabling.

Table 1-4 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 an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the physical interface to be configured, and enter interface configuration mode.
Step 3	speed auto	Configure the interface to autonegotiate speed with the connected device.
Step 4	duplex auto	Configure the interface to autonegotiate duplex mode with the connected device.
Step 5	mdix auto	Enable auto-MDIX on the interface.
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 gigabitethernet0/1
Switch(config-if)# speed auto
Switch(config-if)# duplex auto
Switch(config-if)# mdix auto
Switch(config-if)# end
```

Configuring a Power Management Mode on a PoE Port



Note

PoE commands are supported only when the switch is running the LAN base image. Power over Ethernet Plus (PoE+) is supported only on Catalyst 2960-S switches.

For most situations, the default configuration (auto mode) works well, providing plug-and-play operation. No further configuration is required. However, use the following procedure to give a PoE port higher priority, to make it data only, or to specify a maximum wattage to disallow high-power powered devices on a port.



Note

When you make PoE configuration changes, the port being configured drops power. Depending on the new configuration, the state of the other PoE ports, and the state of the power budget, the port might not be powered up again. For example, port 1 is in the auto and on state, and you configure it for static mode. The switch removes power from port 1, detects the powered device, and repowers the port. If port 1 is in the auto and on state and you configure it with a maximum wattage of 10 W, the switch removes power from the port and then redetects the powered device. The switch repowers the port only if the powered device is a Class 1, Class 2, or a Cisco-only powered device.

Beginning in privileged EXEC mode, follow these steps to configure a power management mode on a PoE-capable port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the physical port to be configured, and enter interface configuration mode.
Step 3	power inline { auto [max <i>max-wattage</i>] never static [max <i>max-wattage</i>]}	<p>Configure the PoE mode on the port. The keywords have these meanings:</p> <ul style="list-style-type: none"> • auto—Enable powered-device detection. If enough power is available, automatically allocate power to the PoE port after device detection. This is the default setting. • (Optional) max <i>max-wattage</i>—Limit the power allowed on the port. The range is 4000 to 15400 milliwatts on a Catalyst 2960 switch and 4000 to 30000 milliwatts on a Catalyst 2960-S switch. If no value is specified, the maximum is allowed. • never—Disable device detection, and disable power to the port. <p>Note If a port has a Cisco powered device connected to it, do not use the power inline never command to configure the port. A false link-up can occur, placing the port into an error-disabled state.</p> <ul style="list-style-type: none"> • static—Enable powered-device detection. Pre-allocate (reserve) power for a port before the switch discovers the powered device. The switch reserves power for this port even when no device is connected and guarantees that power will be provided upon device detection. <p>The switch allocates power to a port configured in static mode before it allocates power to a port configured in auto mode.</p>
Step 4	end	Return to privileged EXEC mode.
Step 5	show power inline [<i>interface-id</i> module <i>switch-number</i>]	The module keyword is applicable only on Catalyst 2960-S switches running the LAN base image.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

For information about the output of the **show power inline** user EXEC command, see the command reference for this release. For more information about PoE-related commands, see the “[Troubleshooting Power over Ethernet Switch Ports](#)” section on page 1-12. For information about configuring voice VLAN, see Chapter 1, “[Configuring Voice VLAN](#).”

Budgeting Power for Devices Connected to a PoE Port

When Cisco powered devices are connected to PoE ports, the switch uses Cisco Discovery Protocol (CDP) to determine the *actual* power consumption of the devices, and the switch adjusts the power budget accordingly. The CDP protocol works with Cisco powered devices and does not apply to IEEE third-party powered devices. For these devices, when the switch grants a power request, the switch adjusts the power budget according to the powered-device IEEE classification. If the powered device is a Class 0 (class status unknown) or a Class 3, the switch budgets 15,400 milliwatts for the device, regardless of

the actual amount of power needed. If the powered device reports a higher class than its actual consumption or does not support power classification (defaults to Class 0), the switch can power fewer devices because it uses the IEEE class information to track the global power budget.

By using the **power inline consumption** *wattage* configuration command, you can override the default power requirement specified by the IEEE classification. The difference between what is mandated by the IEEE classification and what is actually needed by the device is reclaimed into the global power budget for use by additional devices. You can then extend the switch power budget and use it more effectively.

For example, if the switch budgets 15,400 milliwatts on each PoE port, you can connect only 24 Class 0 powered devices. If your Class 0 device power requirement is actually 5000 milliwatts, you can set the consumption *wattage* to 5000 milliwatts and connect up to 48 devices. The total PoE output power available on a 24-port or 48-port switch is 370,000 milliwatts.

**Caution**

You should carefully plan your switch power budget and make certain not to oversubscribe the power supply.

**Note**

When you manually configure the power budget, you must also consider the power loss over the cable between the switch and the powered device.

When you enter the **power inline consumption default** *wattage* or the **no power inline consumption default** global configuration command, or the **power inline consumption** *wattage* or the **no power inline consumption** interface configuration command this caution message appears:

```
%CAUTION: Interface interface-id: Misconfiguring the 'power inline consumption/allocation'
command may cause damage to the switch and void your warranty. Take precaution not to
oversubscribe the power supply.
```

```
It is recommended to enable power policing if the switch supports it.
Refer to documentation.
```

If the power supply is over-subscribed to by up to 20 percent, the switch continues to operate but its reliability is reduced. If the power supply is subscribed to by more than 20 percent, the short-circuit protection circuitry triggers and shuts the switch down.

For more information about the IEEE power classifications, see the [“Power over Ethernet Ports” section on page 1-5](#).

Beginning in privileged EXEC mode, follow these steps to configure the amount of power budgeted to a powered device connected to each PoE port on a switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no cdp run	(Optional) Disable CDP.
Step 3	power inline consumption default <i>wattage</i>	Configure the power consumption of powered devices connected to each the PoE port on the switch. The range for each device is 4000 to 15400 milliwatts on a Catalyst 2960 switch, and 4000 to 30000 milliwatts on a Catalyst 2960-S switch. The default is 15400 milliwatts on a Catalyst 2960 switch, and 30000 milliwatts on a Catalyst 2960-S switch. Note When you use this command, we recommend you also enable power policing.

	Command	Purpose
Step 4	end	Return to privileged EXEC mode.
Step 5	show power inline consumption	Display the power consumption status.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no power inline consumption default** global configuration command.

Beginning in privileged EXEC mode, follow these steps to configure amount of power budgeted to a powered device connected to a specific PoE port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no cdp run	(Optional) Disable CDP.
Step 3	interface <i>interface-id</i>	Specify the physical port to be configured, and enter interface configuration mode.
Step 4	power inline consumption <i>wattage</i>	Configure the power consumption of a powered device connected to a PoE port on the switch. The range for each device is 4000 to 15400 milliwatts on a Catalyst 2960 switch, and 4000 to 30000 milliwatts on a Catalyst 2960-S switch. The default is 15400 milliwatts on a Catalyst 2960 switch, and 30000 milliwatts on a Catalyst 2960-S switch. Note When you use this command, we recommend you also enable power policing.
Step 5	end	Return to privileged EXEC mode.
Step 6	show power inline consumption	Display the power consumption status.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no power inline consumption** interface configuration command.

For information about the output of the **show power inline consumption** privileged EXEC command, see the command reference for this release.

Configuring Power Policing

By default, the switch monitors the real-time power consumption of connected powered devices. You can configure the switch to police the power usage. By default, policing is disabled.

For more information about the cutoff power value, the power consumption values that the switch uses, and the actual power consumption value of the connected device, see the “Power Monitoring and Power Policing” section.

Beginning in privileged EXEC mode, follow these steps to enable policing of the real-time power consumption of a powered device connected to a PoE port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the physical port to be configured, and enter interface configuration mode.
Step 3	power inline police [action log]	<p>If the real-time power consumption exceeds the maximum power allocation on the port, configure the switch to take one of these actions:</p> <ul style="list-style-type: none"> Shut down the PoE port, turn off power to it, and put it in the error-disabled state—Enter the power inline police command. <p>Note You can enable error detection for the PoE error-disabled cause by using the errdisable detect cause inline-power global configuration command. You can also enable the timer to recover from the PoE error-disabled state by using the errdisable recovery cause inline-power interval <i>interval</i> global configuration command.</p> <ul style="list-style-type: none"> Generate a syslog message while still providing power to the port—Enter the power inline police action log command. <p>If you do not enter the action log keywords, the default action shuts down the port and puts the port in the error-disabled state.</p>
Step 4	exit	Return to global configuration mode.
Step 5	errdisable detect cause inline-power and errdisable recovery cause inline-power and errdisable recovery interval <i>interval</i>	<p>(Optional) Enable error recovery from the PoE error-disabled state, and configure the PoE recover mechanism variables.</p> <p>For interval <i>interval</i>, specify the time in seconds to recover from the error-disabled state. The range is 30 to 86400.</p> <p>By default, the recovery interval is 300 seconds.</p>
Step 6	exit	Return to privileged EXEC mode.
Step 7	show power inline police show errdisable recovery	Display the power monitoring status, and verify the error recovery settings.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable policing of the real-time power consumption, use the **no power inline police** interface configuration command. To disable error recovery for PoE error-disabled cause, use the **no errdisable recovery cause inline-power** global configuration command.

For information about the output from the **show power inline police** privileged EXEC command, see the command reference for this release.

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>	Specify the interface for which you are adding a description, and enter interface configuration mode.
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 how to verify the description:

```
Switch# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# description Connects to Marketing
Switch(config-if)# end
Switch# show interfaces gigabitethernet0/2 description
Interface Status      .Protocol Description
Gi0/2    admin down    down      Connects to Marketing
```

Configuring Layer 3 SVIs



Note

Only switches running the LAN base image support Layer 3 SVIs for static routing.

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. You cannot delete VLAN 1.



Note

When you create an SVI, it does not become active until you associate it with a physical port. For information about assigning Layer 2 ports to VLANs, see [Chapter 1, “Configuring VLANs.”](#)

A Layer 3 switch can have an IP address assigned to each SVI, but the switch supports static routing on 16 SVIs. 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.

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

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface vlan <i>vlan-id</i>	Specify the VLAN to be configured as a Layer 3 SVI, and enter interface configuration mode.
Step 3	ip address <i>ip_address subnet_mask</i>	Configure the IP address and IP subnet.
Step 4	end	Return to privileged EXEC mode.
Step 5	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 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

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

This example shows how to configure a Layer 3 SVI and to assign it an IP address:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface vlan 33
Switch(config-if)# ip address 192.20.135.21 255.255.255.0
```

Configuring the System MTU

The default maximum transmission unit (MTU) size for frames received and transmitted on all interfaces is 1500 bytes. You can increase the MTU size for all interfaces operating at 10 or 100 Mb/s 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 mtu jumbo** command. If you do not configure the **system mtu jumbo** command, the setting of the **system mtu** command applies to all Gigabit Ethernet interfaces.

You cannot set the MTU size for an individual interface; you set it for all 10/100 or all Gigabit Ethernet interfaces. When you change the system or jumbo MTU size, you must reset the switch before the new configuration takes effect.

Frames sizes that can be received by the switch CPU are limited to 1998 bytes, no matter what value was entered with the **system mtu** or **system mtu jumbo** commands. Although frames that are forwarded are typically not received by the CPU, in some cases, packets are sent to the CPU, such as traffic sent to control traffic, SNMP, or Telnet.

**Note**

If Layer 2 Gigabit Ethernet interfaces are configured to accept frames greater than the 10/100 interfaces, jumbo frames received on a Layer 2 Gigabit Ethernet interface and sent on a Layer 2 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 Mb/s. The range is 1500 to 1998 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 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 a Gigabit Ethernet port to 1800 bytes:

```
Switch(config)# system mtu jumbo 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 25000
                          ^
% Invalid input detected at '^' marker.
```

Monitoring and Maintaining the Interfaces

These sections contain interface monitoring and maintenance information:

- [Monitoring Interface Status, page 1-38](#)
- [Clearing and Resetting Interfaces and Counters, page 1-38](#)
- [Shutting Down and Restarting the Interface, page 1-39](#)
- [Setting the Type of a Dual-Purpose Uplink Port, page 1-25](#)

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 1-5 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, Release 12.4* from Cisco.com.

Table 1-5 Show Commands for Interfaces

Command	Purpose
show interfaces [<i>interface-id</i>]	(Optional) Display the status and configuration of all interfaces or a specific interface.
show interfaces <i>interface-id</i> status [err-disabled]	(Optional) Display interface status or a list of interfaces in an error-disabled state.
show interfaces [<i>interface-id</i>] switchport	(Optional) Display administrative and operational status of switching ports.
show interfaces [<i>interface-id</i>] description	(Optional) Display the description configured on an interface or all interfaces and the interface status.
show ip interface [<i>interface-id</i>]	(Optional) Display the usability status of all interfaces configured for IP routing or the specified interface.
show interface [<i>interface-id</i>] stats	(Optional) Display the input and output packets by the switching path for the interface.
show interfaces transceiver properties	(Optional) Display speed and duplex settings on the interface.
show interfaces [<i>interface-id</i>] [{ transceiver properties detail }] <i>module number</i>	Display physical and operational status about an SFP module.
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	Display the operational state of the auto-MDIX feature on the interface.
show power inline [<i>interface-id</i>]	Display PoE status for a switch or for an interface.
show power inline police	Display the power policing data.

Clearing and Resetting Interfaces and Counters

Table 1-6 lists the privileged EXEC mode **clear** commands that you can use to clear counters and reset interfaces.

Table 1-6 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 you specify optional arguments that 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.



CHAPTER 1

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 2960 and 2960-S switches](#). It includes information about VLAN membership modes, VLAN configuration modes, VLAN trunks, and dynamic VLAN assignment from a VLAN Membership Policy Server (VMPS). [Unless otherwise noted, the term *switch* refers to a standalone switch and a switch stack.](#)



Note

[Stacking is supported only on Catalyst 2960-S switches running the LAN base image.](#)



Note

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release.

The chapter consists of these sections:

- [Understanding VLANs, page 1-1](#)
- [Configuring Normal-Range VLANs, page 1-4](#)
- [Configuring Extended-Range VLANs, page 1-10](#)
- [Displaying VLANs, page 1-13](#)
- [Configuring VLAN Trunks, page 1-13](#)
- [Configuring VMPS, page 1-22](#)

Understanding VLANs

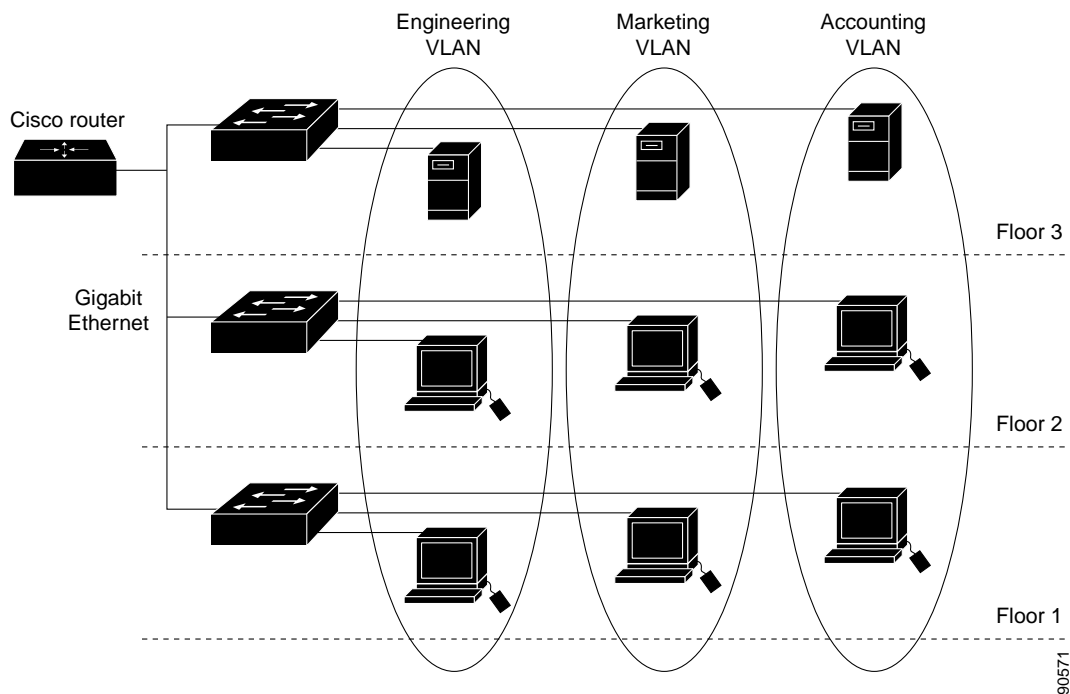
A VLAN is a switched network that is logically segmented by function, project team, or application, without regard to the physical locations of the users. VLANs have the same attributes as physical LANs, but you can group end stations even if they are not physically located on the same LAN segment. Any 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 1-1](#). [VLANs can be formed with ports across the stack.](#) Because a VLAN is considered a separate logical network, it contains its own bridge Management Information Base (MIB) information and can support its own implementation of spanning tree.

**Note**

Before you create VLANs, you must decide whether to use VLAN Trunking Protocol (VTP) to maintain global VLAN configuration for your network.

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

Figure 1-1 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.

Supported VLANs

The switch supports VLANs in VTP client, server, and transparent modes. VLANs are identified by a number from 1 to 4094. VLAN IDs 1002 through 1005 are reserved for Token Ring and FDDI VLANs.

VTP version 1 and version 2 support only normal-range VLANs (VLAN IDs 1 to 1005). In these versions, the switch must be in VTP transparent mode when you create VLAN IDs from 1006 to 4094. Cisco IOS Release 12.2(52)SE and later support VTP version 3. VTP version 3 supports the entire VLAN range (VLANs 1 to 4094). Extended range VLANs (VLANs 1006 to 4094) are supported only in VTP version 3. You cannot convert from VTP version 3 to VTP version 2 if extended VLANs are configured in the domain.

**Note**

Up to 64 VLANs are supported when the switch is running the LAN Lite image.

Although the switch [stack](#) supports a total of [255](#) (normal range and extended range) VLANs, the number of 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 1-6 for more information about the number of spanning-tree instances and the number of VLANs. [The switch supports only IEEE 802.1Q trunking methods for sending VLAN traffic over Ethernet ports.](#)



Note

[Up to 64 spanning-tree instances are supported when the switch is running the LAN Lite image.](#)

VLAN Port Membership Modes

You configure a port to belong to a VLAN by assigning a membership mode that specifies the kind of traffic the port carries and the number of VLANs to which it can belong. [Table 1-1](#) lists the membership modes and membership and VTP characteristics.

Table 1-1 Port Membership Modes and Characteristics

Membership Mode	VLAN Membership Characteristics	VTP Characteristics
Static-access	<p>A static-access port can belong to one VLAN and is manually assigned to that VLAN.</p> <p>For more information, see the “Assigning Static-Access Ports to a VLAN” section on page 1-9.</p>	<p>VTP is not required. If you do not want VTP to globally propagate information, set the VTP mode to transparent. To participate in VTP, there must be at least one trunk port on the switch stack connected to a trunk port of a second switch or switch stack. Stacking is supported only on Catalyst 2960-S switches running the LAN base image.</p>
Trunk (IEEE 802.1Q)	<p>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.</p> <p>For information about configuring trunk ports, see the “Configuring an Ethernet Interface as a Trunk Port” section on page 1-15.</p>	<p>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.</p>

Table 1-1 Port Membership Modes and Characteristics

Membership Mode	VLAN Membership Characteristics	VTP Characteristics
Dynamic access	<p>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 2960 or 2960-S switch. The Catalyst 2960 or 2960-S switch is a VMPS client.</p> <p>You can have dynamic-access ports and trunk ports on the same switch, but you must connect the dynamic-access port to an end station or hub and not to another switch.</p> <p>For configuration information, see the “Configuring Dynamic-Access Ports on VMPS Clients” section on page 1-25.</p>	<p>VTP is required.</p> <p>Configure the VMPS and the client with the same VTP domain name.</p> <p>To participate in VTP, at least one trunk port on the switch stack must be connected to a trunk port of a second switch or switch stack.</p>
Voice VLAN	<p>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.</p>	<p>VTP is not required; it has no effect on a voice VLAN.</p>

For more detailed definitions of access and trunk modes and their functions, see [Table 1-4 on page 1-14](#).

When a port belongs to a VLAN, the switch learns and manages the addresses associated with the port on a per-VLAN basis.

Configuring Normal-Range VLANs

Normal-range VLANs are VLANs with VLAN IDs 1 to 1005. If the switch is in VTP server or VTP 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.)

In VTP versions 1 and 2, the switch must be in VTP transparent mode when you create extended-range VLANs (VLANs with IDs from 1006 to 4094), but these VLANs are not saved in the VLAN database. VTP version 3 supports extended-range VLANs in VTP server and transparent mode. See the “[Configuring Extended-Range VLANs](#)” section on [page 1-10](#).

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 flash memory [on the stack master](#). [Stack members have a vlan.dat file that is consistent with the stack master](#).



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.

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, see the command reference for this release.

These sections contain normal-range VLAN configuration information:

- [Token Ring VLANs, page 1-5](#)
- [Normal-Range VLAN Configuration Guidelines, page 1-6](#)
- [Configuring Normal-Range VLANs, page 1-6](#)
- [Default Ethernet VLAN Configuration, page 1-7](#)
- [Creating or Modifying an Ethernet VLAN, page 1-8](#)
- [Deleting a VLAN, page 1-9](#)
- [Assigning Static-Access Ports to a VLAN, page 1-9](#)

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, see 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 [255](#) 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 the VTP mode is transparent, VTP and VLAN configuration are also saved in the switch running configuration file.
- With VTP versions 1 and 2, the switch supports VLAN IDs 1006 through 4094 only in VTP transparent mode (VTP disabled). These are extended-range VLANs and configuration options are limited. Extended-range VLANs created in VTP transparent mode are not saved in the VLAN database and are not propagated. VTP version 3 supports extended range VLAN (VLANs 1006 to 4094) database propagation. If extended VLANs are configured, you cannot convert from VTP version 3 to version 1 or 2. See the [“Configuring Extended-Range VLANs”](#) section on page 1-10.
- 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 spanning-tree instance.

- [When a switch in a stack learns a new VLAN or deletes or modifies an existing VLAN \(either through VTP over network ports or through the CLI\), the VLAN information is communicated to all stack members.](#)
- [When a switch joins a stack or when stacks merge, VTP information \(the vlan.dat file\) on the new switches will be consistent with the stack master.](#)



Note

[Stacking is supported only on Catalyst 2960-S switches running the LAN base image.](#)

Configuring Normal-Range VLANs

You configure VLANs in `vlan` global configuration command by entering a VLAN ID. Enter a new VLAN ID to create a VLAN, or enter an existing VLAN ID to modify that VLAN. You can use the default VLAN configuration ([Table 1-2](#)) or enter multiple commands to configure the VLAN. For more information about commands available in this mode, see the `vlan` global configuration command

description in the command reference for this release. When you have finished the configuration, you must exit VLAN configuration mode for the configuration to take effect. To display the VLAN configuration, enter the **show vlan** privileged EXEC command.

The configurations of VLAN IDs 1 to 1005 are always saved in the VLAN database (vlan.dat file). If the VTP mode is transparent, they are also saved in the switch running configuration file. You can enter the **copy running-config startup-config** privileged EXEC command to save the configuration in the startup configuration file. [In a switch stack, the whole stack uses the same vlan.dat file and running configuration.](#) 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 selected as follows:

- If the VTP mode is transparent in the startup configuration, and the VLAN database and the VTP domain name from the VLAN database matches that in the startup configuration file, the VLAN database is ignored (cleared), and the VTP and VLAN configurations in the startup configuration file are used. The VLAN database revision number remains unchanged in the VLAN database.
- If the VTP mode or domain name in the startup configuration does not match the VLAN database, the domain name and VTP mode and configuration for the first 1005 VLANs use the VLAN database information.
- In VTP versions 1 and 2, if VTP mode is server, the domain name and VLAN configuration for only the first 1005 VLANs use the VLAN database information. VTP version 3 also supports VLANs 1006 to 4094.

Default Ethernet VLAN Configuration

Table 1-2 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 1-2 Ethernet VLAN Defaults and Ranges

Parameter	Default	Range
VLAN ID	1	1 to 4094. Note Extended-range VLANs (VLAN IDs 1006 to 4094) are only saved in the VLAN database in VTP version 3.
VLAN name	VLANxxxx, where xxxx represents four numeric digits (including leading zeros) equal to the VLAN ID number	No range
IEEE 802.10 SAID	100001 (100000 plus the VLAN ID)	1 to 4294967294
MTU size	1500	1500 to 18190
Translational bridge 1	0	0 to 1005
Translational bridge 2	0	0 to 1005

Table 1-2 Ethernet VLAN Defaults and Ranges (continued)

Parameter	Default	Range
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

With VTP version 1 and 2, if 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 1-10.

For the list of default parameters that are assigned when you add a VLAN, see the “[Configuring Normal-Range VLANs](#)” section on page 1-4.

Beginning in privileged EXEC mode, follow these steps 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 VLAN configuration mode. Enter a new VLAN ID to create a VLAN, or enter an existing VLAN ID to modify that 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 1-10.
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. Note The switch must be running the LAN Base image to use 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** commands.

This example shows how 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
```

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 [stack](#).

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:

	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.

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 1-8.)

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.
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 gigabitethernet0/1
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan 2
Switch(config-if)# end
```

Configuring Extended-Range VLANs

With VTP version 1 and version 2, when the switch is in VTP transparent mode (VTP disabled), you can create extended-range VLANs (in the range 1006 to 4094). VTP version supports extended-range VLANs in server or transparent mode. 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.

With VTP version 1 or 2, 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. Extended-range VLANs created in VTP version 3 are stored in the VLAN database.



Note

Although the switch supports 4094 VLAN IDs, see the “Supported VLANs” section on page 1-2 for the actual number of VLANs supported.

These sections contain extended-range VLAN configuration information:

- [Default VLAN Configuration, page 1-11](#)
- [Extended-Range VLAN Configuration Guidelines, page 1-11](#)
- [Creating an Extended-Range VLAN, page 1-11](#)

Default VLAN Configuration

See [Table 1-2 on page 1-7](#) for the default configuration for Ethernet VLANs. You can change only the MTU size and the remote SPAN configuration state on extended-range VLANs; all other characteristics must remain at the default state.



Note

The switch must be running the LAN Base image to support remote SPAN.

Extended-Range VLAN Configuration Guidelines

Follow these guidelines when creating extended-range VLANs:

- VLAN IDs in the extended range are not saved in the VLAN database and are not recognized by VTP unless the switch is running VTP version 3.
- You cannot include extended-range VLANs in the pruning eligible range.
- In VTP version 1 and 2, a 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. VTP version 3 supports extended VLANs in server and transparent modes.
- For VTP version 1 or 2, you can set the VTP mode to transparent in global configuration mode. 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. If you create extended-range VLANs in VTP version 3, you cannot convert to VTP version 1 or 2.
- 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 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 spanning-tree instance.
- Although the switch [stack](#) supports a total of [255](#) (normal-range and extended-range) VLANs, the number of 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.
- [In a switch stack, the whole stack uses the same running configuration and saved configuration, and extended-range VLAN information is shared across the stack.](#)

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. The extended-range VLAN has the default Ethernet VLAN characteristics (see [Table 1-2](#)) and the MTU size, and RSPAN configuration are the only parameters you can change. See the description of the **vlan** global configuration command in the command reference for the default settings of all parameters. In VTP version 1 or 2, 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 VLAN configuration mode, and the extended-range VLAN is not created.

In VTP version 1 and 2, 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. VTP version 3 saves extended-range VLANs in the VLAN database.

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. Note This step is not required for VTP version 3.
Step 3	vlan <i>vlan-id</i>	Enter an extended-range VLAN ID and enter VLAN configuration 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, 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. <i>RSPAN is supported only if the switch is running the LAN Base image.</i>
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. Note With VTP version 3, the VLAN configuration is also saved in the VLAN database.

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 1-9.

This example shows how to create a new extended-range VLAN with all default characteristics, enter VLAN configuration mode, and save the new VLAN in the switch startup configuration file:

```
Switch(config)# vtp mode transparent
Switch(config)# vlan 2000
Switch(config-vlan)# end
Switch# copy running-config startup config
```

Displaying VLANs

Use the **show vlan** privileged EXEC command to display a list of all VLANs on the switch, including extended-range VLANs. The display includes VLAN status, ports, and configuration information.

[Table 1-3](#) lists the privileged EXEC commands for monitoring VLANs.

Table 1-3 VLAN Monitoring Commands

Command	Purpose
show interfaces [vlan <i>vlan-id</i>]	Display characteristics for all interfaces or for the specified VLAN configured on the switch.
show vlan [id <i>vlan-id</i>]	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, see the command reference for this release.

Configuring VLAN Trunks

These sections contain this conceptual information:

- [Trunking Overview, page 1-13](#)
- [Default Layer 2 Ethernet Interface VLAN Configuration, page 1-15](#)
- [Configuring an Ethernet Interface as a Trunk Port, page 1-15](#)
- [Configuring Trunk Ports for Load Sharing, page 1-19](#)

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. [The switch supports IEEE 802.1Q encapsulation.](#)

You can configure a trunk on a single Ethernet interface or on an EtherChannel bundle.

Ethernet trunk interfaces support different trunking modes (see [Table 1-4](#)). 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.

Table 1-4 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.

IEEE 802.1Q Configuration Considerations

The IEEE 802.1Q trunks impose these limitations on the trunking strategy for a network:

- In a network of Cisco switches connected through IEEE 802.1Q trunks, the switches maintain one spanning-tree instance 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 IEEE 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 IEEE 802.1Q switch. However, spanning-tree information for each VLAN is maintained by Cisco switches separated by a cloud of non-Cisco IEEE 802.1Q switches. The non-Cisco IEEE 802.1Q cloud separating the Cisco switches is treated as a single trunk link between the switches.

- Make sure the native VLAN for an IEEE 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 IEEE 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 IEEE 802.1Q trunk or disable spanning tree on every VLAN in the network. Make sure your network is loop-free before you disable spanning tree.

Default Layer 2 Ethernet Interface VLAN Configuration

Table 1-5 shows the default Layer 2 Ethernet interface VLAN configuration.

Table 1-5 Default Layer 2 Ethernet Interface VLAN Configuration

Feature	Default Setting
Interface mode	switchport mode dynamic auto
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 IEEE 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.

These sections contain this configuration information:

- [Interaction with Other Features, page 1-15](#)
- [Defining the Allowed VLANs on a Trunk, page 1-17](#)
- [Changing the Pruning-Eligible List, page 1-18](#)
- [Configuring the Native VLAN for Untagged Traffic, page 1-18](#)

Interaction with Other Features

Trunking interacts with other features in these ways:

- A trunk port cannot be a secure port.
- Trunk ports can be grouped into EtherChannel port groups, but all trunks in the group must have the same configuration. When a group is first created, all ports follow the parameters set for the first port to be added to the group. If you change the configuration of one of these parameters, the switch propagates the setting 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 IEEE 802.1x on a trunk port, an error message appears, and IEEE 802.1x is not enabled. If you try to change the mode of an IEEE 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 IEEE 802.1x on a dynamic port, an error message appears, and IEEE 802.1x is not enabled. If you try to change the mode of an IEEE 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 a trunk port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port to be configured for trunking, and enter interface configuration mode.
Step 3	switchport mode { dynamic { auto desirable } trunk }	Configure the interface as a Layer 2 trunk (required only if the interface is a Layer 2 access port or to specify the trunking mode). <ul style="list-style-type: none"> • dynamic auto—Set the interface to a trunk link if the neighboring interface is set to trunk or desirable mode. 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 4	switchport access vlan <i>vlan-id</i>	(Optional) Specify the default VLAN, which is used if the interface stops trunking.
Step 5	switchport trunk native vlan <i>vlan-id</i>	Specify the native VLAN for IEEE 802.1Q trunks.
Step 6	end	Return to privileged EXEC mode.
Step 7	show interfaces <i>interface-id</i> switchport	Display the switchport configuration of the interface in the <i>Administrative Mode</i> and the <i>Administrative Trunking Encapsulation</i> fields of the display.
Step 8	show interfaces <i>interface-id</i> trunk	Display the trunk configuration of the interface.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return an interface to its default configuration, use the **default interface** *interface-id* interface configuration command. To reset all trunking characteristics of a trunking interface to the defaults, use the **no switchport trunk** interface configuration command. To disable trunking, use the **switchport mode access** interface configuration command to configure the port as a static-access port.

This example shows how to configure a port as an IEEE 802.1Q trunk. The example assumes that the neighbor interface is configured to support IEEE 802.1Q trunking.

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# switchport mode dynamic desirable
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), DTP, and 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 a trunk:

	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	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, see 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 gigabitethernet0/1
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.

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>	Select the trunk port for which VLANs should be pruned, and enter interface configuration mode.
Step 3	switchport trunk pruning vlan { add except none remove } <i>vlan-list</i> [<i>vlan</i> [, <i>vlan</i> [,..]]]	Configure the list of VLANs allowed to be pruned from the trunk. For explanations about using the add , except , none , and remove keywords, see 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 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 IEEE 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 IEEE 802.1Q configuration issues, see the [“IEEE 802.1Q Configuration Considerations” section on page 1-14](#).

Beginning in privileged EXEC mode, follow these steps to configure the native VLAN on an IEEE 802.1Q trunk:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Define the interface that is configured as the IEEE 802.1Q trunk, and enter interface configuration mode.
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.

Load Sharing Using STP Port Priorities


When two ports on the same switch form a loop, the switch uses the STP port priority to decide 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 1-2 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 1-2 Load Sharing by Using STP Port Priorities



Note

If your switch is a member of a switch stack, you must use the `spanning-tree [vlan vlan-id] cost cost interface configuration command` instead of the `spanning-tree [vlan vlan-id] port-priority priority interface configuration command` to select an interface to put in the forwarding state. Assign lower cost values to interfaces that you want selected first and higher cost values that you want selected last. For more information, see the “Load Sharing Using STP Path Cost” section on page 1-21. Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

Beginning in privileged EXEC mode, follow these steps to configure the network shown in Figure 1-2.

	Command	Purpose
Step 1	<code>configure terminal</code>	Enter global configuration mode on Switch A.
Step 2	<code>vtp domain domain-name</code>	Configure a VTP administrative domain. The domain name can be 1 to 32 characters.
Step 3	<code>vtp mode server</code>	Configure Switch A as the VTP server.
Step 4	<code>end</code>	Return to privileged EXEC mode.
Step 5	<code>show vtp status</code>	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	<code>show vlan</code>	Verify that the VLANs exist in the database on Switch A.
Step 7	<code>configure terminal</code>	Enter global configuration mode.
Step 8	<code>interface interface-id_1</code>	Define the interface to be configured as a trunk, and enter interface configuration mode.
Step 9	<code>switchport mode trunk</code>	Configure the port as a trunk port.
Step 10	<code>end</code>	Return to privileged EXEC mode.
Step 11	<code>show interfaces interface-id_1 switchport</code>	Verify the VLAN configuration.
Step 12		Repeat Steps 7 through 10 on Switch A for a second port in the switch stack.

	Command	Purpose
Step 13		Repeat Steps 7 through 10 on Switch B to configure the trunk ports that connect to the trunk ports configured on Switch A.
Step 14	<code>show vlan</code>	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 15	<code>configure terminal</code>	Enter global configuration mode on Switch A.
Step 16	<code>interface interface-id_1</code>	Define the interface to set the STP port priority, and enter interface configuration mode.
Step 17	<code>spanning-tree vlan 8-10 port-priority 16</code>	Assign the port priority of 16 for VLANs 8 through 10.
Step 18	<code>exit</code>	Return to global configuration mode.
Step 19	<code>interface interface-id_2</code>	Define the interface to set the STP port priority, and enter interface configuration mode.
Step 20	<code>spanning-tree vlan 3-6 port-priority 16</code>	Assign the port priority of 16 for VLANs 3 through 6.
Step 21	<code>end</code>	Return to privileged EXEC mode.
Step 22	<code>show running-config</code>	Verify your entries.
Step 23	<code>copy running-config startup-config</code>	(Optional) Save your entries in the configuration file.

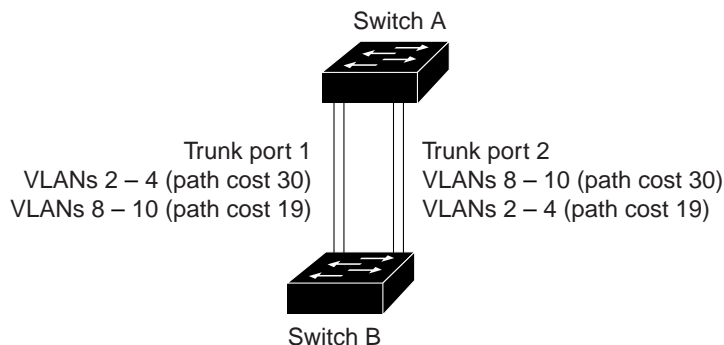
Load Sharing Using STP Path Cost

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 1-3](#), 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 1-3 Load-Sharing Trunks with Traffic Distributed by Path Cost



Beginning in privileged EXEC mode, follow these steps to configure the network shown in [Figure 1-3](#):

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode on Switch A.
Step 2	interface <i>interface-id_1</i>	Define the interface to be configured as a trunk, and enter interface configuration mode.
Step 3	switchport mode trunk	Configure the port as a trunk port.
Step 4	exit	Return to global configuration mode.
Step 5		Repeat Steps 2 through 4 on a second interface in the Switch A stack .
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries. In the display, make sure that the interfaces are configured as trunk ports.
Step 8	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 9	configure terminal	Enter global configuration mode.
Step 10	interface <i>interface-id_1</i>	Define the interface on which to set the STP cost, and enter interface configuration mode.
Step 11	spanning-tree vlan 2-4 cost 30	Set the spanning-tree path cost to 30 for VLANs 2 through 4.
Step 12	end	Return to global configuration mode.
Step 13		Repeat Steps 9 through 12 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 14	exit	Return to privileged EXEC mode.
Step 15	show running-config	Verify your entries. In the display, verify that the path costs are set correctly for both trunk interfaces.
Step 16	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 give 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.

These sections contain this information:

- [“Understanding VMPS” section on page 1-23](#)
- [“Default VMPS Client Configuration” section on page 1-24](#)
- [“VMPS Configuration Guidelines” section on page 1-24](#)
- [“Configuring the VMPS Client” section on page 1-25](#)
- [“Monitoring the VMPS” section on page 1-27](#)

- [“Troubleshooting Dynamic-Access Port VLAN Membership” section on page 1-28](#)
- [“VMPS Configuration Example” section on page 1-28](#)

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 Network Assistant, 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 1-6 shows the default VMPS and dynamic-access port configuration on client switches.

Table 1-6 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.
- IEEE 802.1x ports cannot be configured as dynamic-access ports. If you try to enable IEEE 802.1x on a dynamic-access (VQP) port, an error message appears, and IEEE 802.1x is not enabled. If you try to change an IEEE 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.
- 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	vmpls server ipaddress primary	Enter the IP address of the switch acting as the primary VMPS server.
Step 3	vmpls 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 vmpls	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 **rommand** privileged EXEC command to log in to 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 interface-id	Specify the switch port that is connected to the end station, and enter interface configuration mode.
Step 3	switchport mode access	Set the port to access mode.

	Command	Purpose
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 **recommand** privileged EXEC command to log in to 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	vmmps reconfirm <i>minutes</i>	Enter the number of minutes between reconfirmations of the dynamic VLAN membership. The range is 1 to 120. The default is 60 minutes.
Step 3	end	Return to privileged EXEC mode.
Step 4	show vmmps	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 vmmps 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	vmpls retry count	Change the retry count. The retry range is 1 to 10; the default is 3.
Step 3	end	Return to privileged EXEC mode.
Step 4	show vmpls	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 vmpls retry** global configuration command.

Monitoring the VMPS

You can display information about the VMPS by using the **show vmpls** 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 expires, or you can force it by entering the **vmpls reconfirm** privileged EXEC command or its Network Assistant or SNMP equivalent.

This is an example of output for the **show vmpls** privileged EXEC command:

```
Switch# show vmpls
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 1-4](#) 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 1-4 Dynamic Port VLAN Membership Configuration





CHAPTER 1

Configuring VTP

This chapter describes how to use the VLAN Trunking Protocol (VTP) and the VLAN database for managing VLANs with the Catalyst 2960 and 2960-S switches. Unless otherwise noted, the term *switch* refers to a standalone switch and a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release.

The chapter consists of these sections:

- [Understanding VTP, page 1-1](#)
- [Configuring VTP, page 1-9](#)
- [Monitoring VTP, page 1-18](#)

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.

VTP functionality is supported across the stack, and all switches in the stack maintain the same VLAN and VTP configuration inherited from the stack master. When a switch learns of a new VLAN through VTP messages or when a new VLAN is configured by the user, the new VLAN information is communicated to all switches in the stack.

When a switch joins the stack or when stacks merge, the new switches get VTP information from the stack master.

The switch supports 255 VLANs, but the number of 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.

**Note**

The switch supports up to 64 VLANs when it is running the LAN Lite image.

VTP version 1 and version 2 support only normal-range VLANs (VLAN IDs 1 to 1005). Cisco IOS Release 12.2(52)SE and later support VTP version 3. VTP version 3 supports the entire VLAN range (VLANs 1 to 4094). Extended range VLANs (VLANs 1006 to 4094) are supported only in VTP version 3. You cannot convert from VTP version 3 to VTP version 2 if extended VLANs are configured in the domain.

**Note**

The switch must be running the LAN base image to support VTP version 3.

These sections contain this conceptual information:

- [The VTP Domain, page 1-2](#)
- [VTP Modes, page 1-3](#)
- [VTP Advertisements, page 1-4](#)
- [VTP Version 2, page 1-5](#)
- [VTP Version 3, page 1-5](#)
- [VTP Pruning, page 1-6](#)

The VTP Domain

A VTP domain (also called a VLAN management domain) consists of one switch or several interconnected switches or switch stacks 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 default, the switch is in the 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 1-17](#) 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 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 1-10.

VTP Modes

You can configure a supported switch stack to be in one of the VTP modes listed in [Table 1-1](#).

Table 1-1 VTP Modes

VTP Mode	Description
VTP server	<p>In VTP server mode, you can create, modify, and delete VLANs, and specify other configuration parameters (such as the VTP version) for the entire VTP domain. VTP servers advertise their VLAN configurations to other switches in the same VTP domain and synchronize their VLAN configurations with other switches based on advertisements received over trunk links.</p> <p>VTP server is the default mode.</p> <p>Note In VTP server mode, VLAN configurations are saved in NVRAM. If the switch detects a failure while writing a configuration to NVRAM, VTP mode automatically changes from server mode to client mode. If this happens, the switch cannot be returned to VTP server mode until the NVRAM is functioning.</p>
VTP client	<p>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.</p> <p>In VTP versions 1 and 2, in VTP client mode, VLAN configurations are not saved in NVRAM. In VTP version 3, VLAN configurations are saved in NVRAM in client mode.</p>

Table 1-1 VTP Modes (continued)

VTP Mode	Description
VTP transparent	<p>VTP transparent switches do not participate in VTP. A VTP transparent switch does not advertise its VLAN configuration and does not synchronize its VLAN configuration based on received advertisements. However, in VTP version 2 or version 3, transparent switches do forward VTP advertisements that they receive from other switches through their trunk interfaces. You can create, modify, and delete VLANs on a switch in VTP transparent mode.</p> <p>In VTP versions 1 and 2, the switch must be in VTP transparent mode when you create extended-range VLANs. VTP version 3 also supports creating extended-range VLANs in client or server mode. See the “Configuring Extended-Range VLANs” section on page 1-10.</p> <p>When the switch is in VTP transparent mode, the VTP and VLAN configurations are saved in NVRAM, but they are not advertised to other switches. In this mode, VTP mode and domain name are saved in the switch running configuration, and you can save this information in the switch startup configuration file by using the copy running-config startup-config privileged EXEC command. The running configuration and the saved configuration are the same for all switches in a stack.</p>
VTP off	A switch in VTP off mode functions in the same manner as a VTP transparent switch, except that it does not forward VTP advertisements on trunks.

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 stack 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 1-13.

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 (IEEE 802.1Q)
- VLAN name
- VLAN type
- VLAN state
- Additional VLAN configuration information specific to the VLAN type

In VTP version 3, VTP advertisements also include the primary server ID, an instance number, and a start index.

VTP Version 2

If you use VTP in your network, you must decide which version of VTP to use. By default, VTP operates in version 1.

VTP version 2 supports these features that are 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 1-4.
- 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. Although VTP version 2 supports only one domain, a VTP version 2 transparent switch forwards a message only when the domain name matches.
- 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 Version 3

VTP version 3 supports these features that are not supported in version 1 or version 2:

- Enhanced authentication—You can configure the authentication as **hidden** or **secret**. When **hidden**, the secret key from the password string is saved in the VLAN database file, but it does not appear in plain text in the configuration. Instead, the key associated with the password is saved in hexadecimal format in the running configuration. You must reenter the password if you enter a takeover command in the domain. When you enter the **secret** keyword, you can directly configure the password secret key.
- Support for extended range VLAN (VLANs 1006 to 4094) database propagation. VTP versions 1 and 2 propagate only VLANs 1 to 1005. If extended VLANs are configured, you cannot convert from VTP version 3 to version 1 or 2.



Note VTP pruning still applies only to VLANs 1 to 1005, and VLANs 1002 to 1005 are still reserved and cannot be modified.

- Private VLAN support.
- Support for any database in a domain. In addition to propagating VTP information, version 3 can propagate Multiple Spanning Tree (MST) protocol database information. A separate instance of the VTP protocol runs for each application that uses VTP.

- VTP primary server and VTP secondary servers. A VTP primary server updates the database information and sends updates that are honored by all devices in the system. A VTP secondary server can only back up the updated VTP configurations received from the primary server to its NVRAM.

By default, all devices come up as secondary servers. You can enter the **vtp primary** privileged EXEC command to specify a primary server. Primary server status is only needed for database updates when the administrator issues a takeover message in the domain. You can have a working VTP domain without any primary servers. Primary server status is lost if the device reloads or domain parameters change, even when a password is configured on the switch.

- The option to turn VTP on or off on a per-trunk (per-port) basis. You can enable or disable VTP per port by entering the **[no] vtp** interface configuration command. When you disable VTP on trunking ports, all VTP instances for that port are disabled. You cannot set VTP to *off* for the MST database and *on* for the VLAN database on the same port.

When you globally set VTP mode to off, it applies to all the trunking ports in the system. However, you can specify on or off on a per-VTP instance basis. For example, you can configure the switch as a VTP server for the VLAN database but with VTP *off* for the MST database.

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 in all VTP versions.

Figure 1-1 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 1-1 Flooding Traffic without VTP Pruning


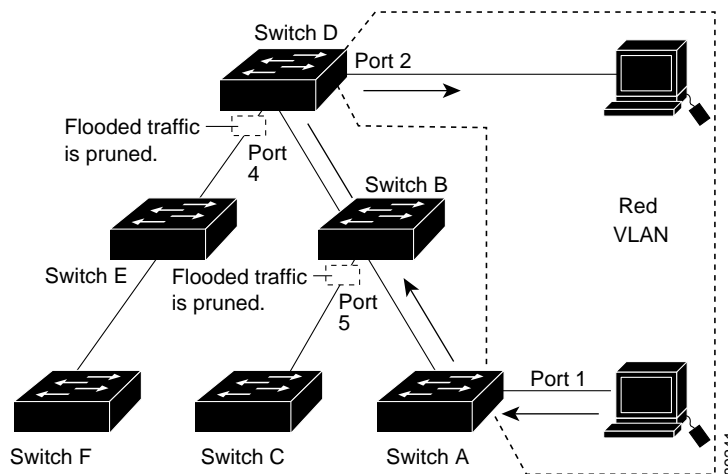


Figure 1-2 shows a switched network with VTP pruning enabled. The broadcast traffic from Switch A is not forwarded to Switches C, E, and F because traffic for the Red VLAN has been pruned on the links shown (Port 5 on Switch B and Port 4 on Switch D).

Figure 1-2 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 1-16. 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 1-18). 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.

VTP and Switch Stacks

**Note**

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

VTP configuration is the same in all members of a switch stack. When the switch stack is in VTP server or client mode, all switches in the stack carry the same VTP configuration. When VTP mode is transparent or off, the stack is not taking part in VTP.

- When a switch joins the stack, it inherits the VTP and VLAN properties of the stack master.
- All VTP updates are carried across the stack.
- When VTP mode is changed in a switch in the stack, the other switches in the stack also change VTP mode, and the switch VLAN database remains consistent.

VTP version 3 functions the same on a standalone switch or a stack except when the switch stack is the primary server for the VTP database. In this case, the MAC address of the stack master is used as the primary server ID. If the master switch reloads or is powered off, a new stack master is elected.

- If you do not configure the persistent MAC address feature (by entering the **stack-mac persistent timer [0 | time-value]** global configuration command, when the new master is elected, it sends a takeover message with the new master MAC address as the primary server.
- If persistent MAC address is configured, the new master waits for the configured **stack-mac persistent timer** value. If the previous master switch does not rejoin the stack during this time, then the new master issues the takeover message.

For more information about the switch stack, see [Chapter 1, “Managing Switch Stacks.”](#)

Configuring VTP

These sections contain this configuration information:

- [Default VTP Configuration, page 1-10](#)
- [VTP Configuration Guidelines, page 1-10](#)
- [Configuring VTP Mode, page 1-12](#)
- [Enabling the VTP Version, page 1-15](#)
- [Enabling VTP Pruning, page 1-16](#)
- [Configuring VTP on a Per-Port Basis, page 1-17](#)
- [Adding a VTP Client Switch to a VTP Domain, page 1-17](#)

Default VTP Configuration

Table 1-2 shows the default VTP configuration.

Table 1-2 Default VTP Configuration

Feature	Default Setting
VTP domain name	Null.
VTP mode (VTP version 1 and version 2)	Server.
VTP mode (VTP version 3)	The mode is the same as the mode in VTP version 1 or 2 before conversion to version 3.
VTP version	Version 1.
MST database mode	Transparent.
VTP version 3 server type	Secondary.
VTP password	None.
VTP pruning	Disabled.

VTP Configuration Guidelines

You 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, see 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 if the switch resets.

When you save VTP information in the switch startup configuration file and restart the switch, the configuration is selected as follows:

- If the VTP mode is transparent in both 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). 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 the domain name in the startup configuration do not match the VLAN database, the domain name and the VTP mode and configuration for the first 255 VLANs use the VLAN database information.

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 have the same domain name, but they do not need to 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).
- If a switch running VTP version 1 but capable of running VTP version 2 receives VTP version 3 advertisements, it automatically moves to VTP version 2.
- If a switch running VTP version 3 is connected to a switch running VTP version 1, the VTP version 1 switch moves to VTP version 2, and the VTP version 3 switch sends scaled-down versions of the VTP packets so that the VTP version 2 switch can update its database.
- A switch running VTP version 3 cannot move to version 1 or 2 if it has extended VLANs.
- 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 that have version 2 enabled.
- We recommend placing VTP version 1 and 2 switches at the edge of the network because they do not forward VTP version 3 advertisements.

- If there are TrBRF and TrCRF Token Ring networks in your environment, you must enable VTP version 2 or version 3 for Token Ring VLAN switching to function properly. To run Token Ring and Token Ring-Net, disable VTP version 2.
- VTP version 1 and version 2 do not propagate configuration information for extended range VLANs (VLANs 1006 to 4094). You must configure these VLANs manually on each device. VTP version 3 supports extended-range VLANs. You cannot convert from VTP version 3 to VTP version 2 if extended VLANs are configured.
- When a VTP version 3 device trunk port receives messages from a VTP version 2 device, it sends a scaled-down version of the VLAN database on that particular trunk in VTP version 2 format. A VTP version 3 device does not send VTP version 2-formatted packets on a trunk unless it first receives VTP version 2 packets on that trunk port.
- When a VTP version 3 device detects a VTP version 2 device on a trunk port, it continues to send VTP version 3 packets, in addition to VTP version 2 packets, to allow both kinds of neighbors to coexist on the same trunk.
- A VTP version 3 device does not accept configuration information from a VTP version 2 or version 1 device.
- Two VTP version 3 regions can only communicate in transparent mode over a VTP version 1 or version 2 region.
- Devices that are only VTP version 1 capable cannot interoperate with VTP version 3 devices.

Configuration Requirements

When you configure VTP, you must configure a trunk port on the switch stack 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 1-13](#).

If you are configuring VTP on a cluster member switch to a VLAN, use the **rcommand** privileged EXEC command to log in to the member switch. For more information about the command, see the command reference for this release.

In VTP versions 1 and 2, when you configure extended-range VLANs on the switch, the switch must be in VTP transparent mode. VTP version 3 also supports creating extended-range VLANs in client or server mode.

Configuring VTP Mode

You can configure VTP mode as one of these:

- When a switch is in VTP server mode, you can change the VLAN configuration and have it propagated throughout the network.
- 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.
- When you configure the switch for VTP transparent mode, VTP is disabled 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 its trunk links.
- VTP off mode is the same as VTP transparent mode except that VTP advertisements are not forwarded.

Follow these guidelines:

- For VTP version 1 and version 2, if extended-range VLANs are configured on the switch, you cannot change VTP mode to client or server. You receive an error message, and the configuration is not allowed. VTP version 1 and version 2 do not propagate configuration information for extended range VLANs (VLANs 1006 to 4094). You must manually configure these VLANs on each device.



Note For VTP version 1 and 2, 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 starts 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).

- VTP version 3 supports extended-range VLANs. If extended VLANs are configured, you cannot convert from VTP version 3 to VTP version 2.
- If you configure the switch for VTP client mode, the switch does not create the VLAN database file (vlan.dat). If the switch is then powered off, it resets the VTP configuration to the default. To keep the VTP configuration with VTP client mode after the switch restarts, you must first configure the VTP domain name before the VTP mode.



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 VTP mode:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vtp domain <i>domain-name</i>	Configure the VTP administrative-domain name. The name can be 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. This command is optional for modes other than server mode. VTP server mode requires a domain name. If the switch has a trunk connection to a VTP domain, the switch learns the domain name from the VTP server in the domain. You should configure the VTP domain before configuring other VTP parameters.
Step 3	vtp mode { client server transparent off } { vlan mst unknown }	Configure the switch for VTP mode (client, server, transparent or off). (Optional) Configure the database: <ul style="list-style-type: none"> vlan—the VLAN database is the default if none are configured. mst—the multiple spanning tree (MST) database. unknown—an unknown database type.

	Command	Purpose
Step 4	vtp password <i>password</i>	(Optional) Set the password for the VTP domain. The password can be 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. See the “ Configuring a VTP Version 3 Password ” section on page 1-14 for options available with VTP version 3.
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.
Step 7	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.

When you configure a domain name, it cannot be removed; you can only reassign a switch to a different domain. To return a switch in another mode to VTP server mode, use the **no vtp mode** global configuration command. To return the switch to a no-password state, use the **no vtp password** global configuration command.

This example shows how to configure the switch as a VTP server with the domain name *eng_group* and the password *mypassword*:

```
Switch(config)# vtp domain eng_group
Setting VTP domain name to eng_group.
Switch(config)# vtp mode server
Setting device to VTP Server mode for VLANs.
Switch(config)# vtp password mypassword
Setting device VLAN database password to mypassword.
Switch(config)# end
```

Configuring a VTP Version 3 Password

Beginning in privileged EXEC mode, follow these steps to configure the password when using VTP version 3:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vtp password <i>password</i> [hidden secret]	(Optional) Set the password for the VTP domain. The password can be 8 to 64 characters. <ul style="list-style-type: none"> (Optional) hidden—Enter hidden to ensure that the secret key generated from the password string is saved in the nvam:vlan.dat file. If you configure a takeover by configuring a VTP primary server, you are prompted to reenter the password. (Optional) secret—Enter secret to directly configure the password. The secret password must contain 32 hexadecimal characters.
Step 3	end	Return to privileged EXEC mode.

	Command	Purpose
Step 4	show vtp password	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save the configuration in the startup configuration file.

To clear the password, enter the **no vtp password** global configuration command.

This example shows how to configure a hidden password and how it appears.

```
Switch(config)# vtp password mypassword hidden
Generating the secret associated to the password.
Switch(config)# end
Switch# show vtp password
VTP password: 89914640C8D90868B6A0D8103847A733
```

Configuring a VTP Version 3 Primary Server

Beginning in privileged EXEC mode, follow these steps on a VTP server to configure it as a VTP primary server (version 3 only), which starts a takeover operation:

	Command	Purpose
Step 1	vtp primary-server [vlan mst] [force]	<p>Change the operational state of a switch from a secondary server (the default) to a primary server and advertise the configuration to the domain. If the switch password is configured as hidden, you are prompted to reenter the password.</p> <ul style="list-style-type: none"> (Optional) vlan—Select the VLAN database as the takeover feature. This is the default. (Optional) mst—Select the multiple spanning tree (MST) database as the takeover feature. (Optional) force—Entering force overwrites the configuration of any conflicting servers. If you do not enter force, you are prompted for confirmation before the takeover.

This example shows how to configure a switch as the primary server for the VLAN database (the default) when a hidden or secret password was configured:

```
Switch# vtp primary vlan
Enter VTP password: mypassword
This switch is becoming Primary server for vlan feature in the VTP domain

VTP Database Conf Switch ID      Primary Server Revision System Name
-----
VLANDB          Yes  00d0.00b8.1400=00d0.00b8.1400 1          stp7

Do you want to continue (y/n) [n]? y
```

Enabling the VTP Version

VTP version 2 and version 3 are disabled by default.



Note

The switch must be running the LAN base image to support VTP version 3.

- When you enable VTP version 2 on a switch, every VTP version 2-capable switch in the VTP domain enables version 2. To enable VTP version 3, you must manually configure it on each switch.
- With VTP versions 1 and 2, you can configure the version only on switches in VTP server or transparent mode. If a switch is running VTP version 3, you can change to version 2 when the switch is in client mode if no extended VLANs exist, no private VLANs exist, and no hidden password was configured.

**Caution**

VTP version 1 and VTP version 2 are not interoperable on switches in the same VTP domain. Do not enable VTP version 2 unless every switch in the VTP domain supports version 2.

- In TrCRF and TrBRF Token ring environments, you must enable VTP version 2 or VTP version 3 for Token Ring VLAN switching to function properly. For Token Ring and Token Ring-Net media, disable VTP version 2 must be disabled.
- VTP version 3 is supported on switches running Cisco IOS Release 12.2(52) SE or later.

**Caution**

In VTP version 3, both the primary and secondary servers can exist on an instance in the domain.

For more information on VTP version configuration guidelines, see the [“VTP Version” section on page 1-11](#).

Beginning in privileged EXEC mode, follow these steps to configure the VTP version:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vtp version {1 2 3}	Enable the VTP version on the switch. The default is VTP version 1.
Step 3	end	Return to privileged EXEC mode.
Step 4	show vtp status	Verify that the configured VTP version is enabled.
Step 5	copy running-config startup-config	(Optional) Save the configuration in the startup configuration file.

To return to the default VTP version 1, use the **no vtp version** global 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.

	Command	Purpose
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.

With VTP versions 1 and 2, when you enable pruning on the VTP server, it is enabled for the entire VTP domain. In VTP version 3, you must manually enable pruning on each switch in the 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 1-18.

Configuring VTP on a Per-Port Basis

With VTP version 3, you can enable or disable VTP on a per-port basis. You can enable VTP only on ports that are in trunk mode. Incoming and outgoing VTP traffic are blocked, not forwarded.

Beginning in privileged EXEC mode, follow these steps to enable VTP on a port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Identify an interface, and enter interface configuration mode.
Step 3	vtp	Enable VTP on the specified port.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config interface <i>interface-id</i>	Verify the change to the port.
Step 6	show vtp status	Verify the configuration.

To disable VTP on the interface, use the **no vtp** interface configuration command.

```
Switch(config)# interface gigabitethernet 0/1
Switch(config-if)# vtp
Switch(config-if)# end
```

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. With VTP versions 1 and 2, adding a switch that has a revision number higher than the revision number in the VTP domain can erase all VLAN information from the VTP server and VTP domain. With VTP version 3, the VLAN information is not erased.

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: <ol style="list-style-type: none"> 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.

After resetting the configuration revision number, add the switch to the VTP domain.



Note

You can use the **vtp mode transparent** global configuration command to disable VTP on the switch and then to 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 1-3 shows the privileged EXEC commands for monitoring VTP activity.

Table 1-3 VTP Monitoring Commands

Command	Purpose
show vtp counters	Display counters about VTP messages that have been sent and received.
show vtp devices [conflict]	Display information about all VTP version 3 devices in the domain. Conflicts are VTP version 3 devices with conflicting primary servers. The show vtp devices command does not display information when the switch is in transparent or off mode.

Table 1-3 VTP Monitoring Commands (continued)

Command	Purpose
show vtp interface [<i>interface-id</i>]	Display VTP status and configuration for all interfaces or the specified interface.
show vtp password	Display the VTP password. The form of the password displayed depends on whether or not the hidden keyword was entered and if encryption is enabled on the switch.
show vtp status	Display the VTP switch configuration information.



CHAPTER 1

Configuring Voice VLAN

This chapter describes how to configure the voice VLAN feature on the Catalyst 2960 and 2960-S switches. Unless otherwise noted, the term *switch* refers to a standalone switch and a switch stack. Voice VLAN is referred to as an *auxiliary VLAN* in some Catalyst 6500 family switch documentation.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.



Note

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release.

This chapter consists of these sections:

- [Understanding Voice VLAN, page 1-1](#)
- [Configuring Voice VLAN, page 1-3](#)
- [Displaying Voice VLAN, page 1-7](#)

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 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 1, “Configuring QoS.”](#)


The Cisco 7960 IP Phone is a configurable device, and you can configure it to forward traffic with an IEEE 802.1p priority. You can configure the switch to trust or override the traffic priority assigned by a Cisco IP Phone.

The Cisco IP Phone contains an integrated three-port 10/100 switch as shown in [Figure 1-1](#). 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 1-1 shows one way to connect a Cisco 7960 IP Phone.

Figure 1-1 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 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 IEEE 802.1Q or IEEE 802.1p frame types) from the device attached to the access port on the Cisco IP Phone (see Figure 1-1). You can configure Layer 2 access ports on the switch to send CDP packets that instruct the attached phone to configure the 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 phone unchanged.
- In untrusted mode, all traffic in IEEE 802.1Q or IEEE 802.1p frames received through the access port on the Cisco 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 phone unchanged, regardless of the trust state of the access port on the phone.

Configuring Voice VLAN

- [Default Voice VLAN Configuration, page 1-3](#)
- [Voice VLAN Configuration Guidelines, page 1-3](#)
- [Configuring a Port Connected to a Cisco 7960 IP Phone, page 1-4](#)

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 IEEE 802.1p or IEEE 802.1Q tagged traffic.

Voice VLAN Configuration Guidelines

These are the voice VLAN configuration guidelines:

- Voice VLAN configuration is only supported on switch access ports; voice VLAN configuration is not supported on trunk ports.

**Note**

Trunk ports can carry any number of voice VLANs, similar to regular VLANs. The configuration of voice VLANs is not required on trunk ports.

- The voice VLAN should be present and active on the switch for the IP phone to correctly communicate on the voice VLAN. Use the **show vlan** privileged EXEC command to see if the VLAN is present (listed in the display). If the VLAN is not listed, see [Chapter 1, “Configuring VLANs,”](#) for information on how to create the voice VLAN.
- The Power over Ethernet (PoE) switches are capable of automatically providing power to Cisco pre-standard and IEEE 802.3af-compliant powered devices if they are not being powered by an AC power source. For information about PoE interfaces, see the [“Configuring a Power Management Mode on a PoE Port”](#) section on page 1-30.
- 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 1, “Configuring QoS.”](#)
- You must enable CDP on the switch port connected to the Cisco IP Phone to send the configuration to the phone. (CDP is globally enabled by default 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 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 IEEE 802.1p or untagged frames.
 - The Cisco IP Phone uses IEEE 802.1p frames, and the device uses untagged frames.
 - The Cisco IP Phone uses untagged frames, and the device uses IEEE 802.1p frames.
 - The Cisco IP Phone uses IEEE 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 1-25](#) for more information.
 - IEEE 802.1x authenticated port. See the [“Configuring 802.1x Readiness Check” section on page 1-38](#) for more information.



Note If you enable IEEE 802.1x on an access port on which a voice VLAN is configured and to which a Cisco IP Phone is connected, the phone loses connectivity to the switch for up to 30 seconds.

- Protected port. See the [“Configuring Protected Ports” section on page 1-6](#) for more information.
- A source or destination port for a SPAN or RSPAN session.
- Secure port. See the [“Configuring Port Security” section on page 1-8](#) 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 phone requires up to two MAC addresses. The phone address is learned on the voice VLAN and might also be learned on the access VLAN. Connecting a PC to the 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 decide how the Cisco IP Phone carries voice traffic and data traffic.

These sections contain this configuration information:

- [Configuring Cisco IP Phone Voice Traffic, page 1-5](#)
- [Configuring the Priority of Incoming Data Frames, page 1-6](#)

Configuring Cisco 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 IEEE 802.1Q frames for a specified voice VLAN with a Layer 2 CoS value. It can use IEEE 802.1p priority tagging to give voice traffic a higher priority and forward all voice traffic through the native (access) VLAN. The Cisco 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>	Specify the interface connected to the phone, and enter interface configuration mode.
Step 3	mls qos trust cos	Configure the interface to classify incoming 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.
Step 4	switchport voice vlan { <i>vlan-id</i> dot1p none untagged }	Configure how the Cisco IP Phone carries voice traffic: <ul style="list-style-type: none"> vlan-id—Configure the phone to forward all voice traffic through the specified VLAN. By default, the Cisco IP Phone forwards the voice traffic with an IEEE 802.1Q priority of 5. Valid VLAN IDs are 1 to 4094. dot1p—Configure the switch to accept voice and data IEEE 802.1p priority frames tagged with VLAN ID 0 (the native VLAN). By default, the switch drops all voice and data traffic tagged with VLAN 0. If configured for 802.1p the Cisco IP Phone forwards the traffic with an IEEE 802.1p priority of 5. none—Allow the 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 a Cisco IP Phone to use the CoS value to classify incoming traffic and to accept voice and data priority traffic tagged with VLAN ID 0:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet0/1
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



Note

To set priority of incoming data frames, the switch must be running the LAN Base image.

You can connect a PC or other data device to a Cisco IP Phone port. To process tagged data traffic (in IEEE 802.1Q or IEEE 802.1p frames), you can configure the switch to send CDP packets to instruct the 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 phone to not change (trust) or to override (not trust) the priority of frames arriving on the 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>	Specify the interface connected to the Cisco IP Phone, and enter interface configuration mode.
Step 3	switchport priority extend { <i>cos value</i> trust }	Set the priority of data traffic received from the Cisco IP Phone access port: <ul style="list-style-type: none"> cos value—Configure the 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 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 a Cisco 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 gigabitethernet0/1
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.



CHAPTER 1

Configuring STP

This chapter describes how to configure the Spanning Tree Protocol (STP) on port-based VLANs on the Catalyst 2960 and 2960-S switches. The switch can use either the per-VLAN spanning-tree plus (PVST+) protocol based on the IEEE 802.1D standard and Cisco proprietary extensions, or the rapid per-VLAN spanning-tree plus (rapid-PVST+) protocol based on the IEEE 802.1w standard. A switch stack appears as a single spanning-tree node to the rest of the network, and all stack members use the same bridge ID. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

For information about the Multiple Spanning Tree Protocol (MSTP) and how to map multiple VLANs to the same spanning-tree instance, see [Chapter 1, “Configuring MSTP.”](#) For information about other spanning-tree features such as Port Fast, UplinkFast, root guard, and so forth, see [Chapter 1, “Configuring Optional Spanning-Tree Features.”](#)



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

This chapter consists of these sections:

- [Understanding Spanning-Tree Features, page 1-1](#)
- [Configuring Spanning-Tree Features, page 1-12](#)
- [Displaying the Spanning-Tree Status, page 1-24](#)

Understanding Spanning-Tree Features

These sections contain this conceptual information:

- [STP Overview, page 1-2](#)
- [Spanning-Tree Topology and BPDUs, page 1-3](#)
- [Bridge ID, Switch Priority, and Extended System ID, page 1-4](#)
- [Spanning-Tree Interface States, page 1-5](#)
- [How a Switch or Port Becomes the Root Switch or Root Port, page 1-8](#)
- [Spanning Tree and Redundant Connectivity, page 1-8](#)

- [Spanning-Tree Address Management, page 1-9](#)
- [Accelerated Aging to Retain Connectivity, page 1-9](#)
- [Spanning-Tree Modes and Protocols, page 1-10](#)
- [Supported Spanning-Tree Instances, page 1-10](#)
- [Spanning-Tree Interoperability and Backward Compatibility, page 1-11](#)
- [STP and IEEE 802.1Q Trunks, page 1-11](#)

For configuration information, see the “Configuring Spanning-Tree Features” section on page 1-12.

For information about optional spanning-tree features, see Chapter 1, “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 bridge in the spanning tree
- **Backup**—A blocked port in a loopback configuration

The switch that has *all* of its ports as the designated role or as the backup role is the root switch. The switch that has at least *one* of its ports in the designated role is called the designated switch.

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.



Note

The default is for the switch to send keepalive messages (to detect local loopback conditions) only on interfaces that do not have small form-factor pluggable (SFP) modules. You can use the **[no] keepalive** interface configuration command to change the default for an interface.

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. In a switch stack, all switches use the same bridge ID for a given spanning-tree instance.
- 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). In a switch stack, one stack member is elected as the stack root switch. The stack root switch contains the outgoing root port (Switch 1), as shown in [Figure 1-1 on page 1-4](#).

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 1-1 on page 1-5](#).

- 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.

When selecting the root port on a switch stack, spanning tree follows this sequence:


- Selects the lowest root bridge ID
- Selects the lowest path cost to the root switch
- Selects the lowest designated bridge ID
- Selects the lowest designated path cost
- Selects the lowest port ID

Only one outgoing port on the stack root switch is selected as the root port. The remaining switches in the stack become its designated switches (Switch 2 and Switch 3) as shown in [Figure 1-1 on page 1-4](#).

- 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.

Figure 1-1 Spanning-Tree Port States in a Switch Stack

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 a different bridge IDs for each configured VLAN. Each VLAN on the switch has a unique 8-byte bridge ID. The 2 most-significant bytes are used for the switch priority, and the remaining 6 bytes are derived from the switch MAC address.

The switch supports the IEEE 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 1-1](#), the 2 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 1-1 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. Because the switch stack appears as a single switch to the rest of the network, all switches in the stack use the same bridge ID for a given spanning tree. If the stack master fails, the stack members recalculate their bridge IDs of all running spanning trees based on the new MAC address of the new stack master.

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 1-16, the [“Configuring a Secondary Root Switch”](#) section on page 1-18, and the [“Configuring the Switch Priority of a VLAN”](#) section on page 1-21.

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 decides 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.

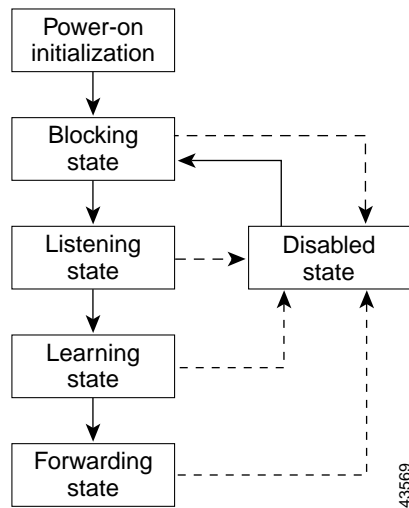
An interface moves through these states:

- 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 1-2 illustrates how an interface moves through the states.

Figure 1-2 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 move 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 decides 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 1-3](#), 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 1-3 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 1-4](#). 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 1-4 Spanning Tree and Redundant Connectivity



You can also create redundant links between switches by using EtherChannel groups. For more information, see [Chapter 1, “Configuring EtherChannels and Link-State Tracking.”](#)

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 in the stack receives but does not forward packets destined for addresses between 0x0180C2000000 and 0x0180C200000F.

If spanning tree is enabled, the CPU on each switch in the stack receives packets destined for 0x0180C2000000 and 0x0180C2000010. If spanning tree is disabled, each switch in the stack 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 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 re-provision your network. In rapid-PVST+ mode, each VLAN runs its own spanning-tree instance up to the maximum supported.

- **MSTP**—This spanning-tree mode is based on the IEEE 802.1s standard. You can map multiple VLANs to the same spanning-tree instance, which reduces the number of spanning-tree instances required to support a large number of VLANs. The MSTP runs on top of the RSTP (based on IEEE 802.1w), which provides for rapid convergence of the spanning tree by eliminating the forward delay and by quickly transitioning root ports and designated ports to the forwarding state. In a switch stack, the cross-stack rapid transition (CSRT) feature performs the same function as RSTP. You cannot run MSTP without RSTP or CSRT.

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 1, “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 stack supports up to 128 spanning-tree instances.

In MSTP mode, the switch stack supports up to 65 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 1-13.

Spanning-Tree Interoperability and Backward Compatibility

Table 1-2 lists the interoperability and compatibility among the supported spanning-tree modes in a network.

Table 1-2 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.

All stack members run the same version of spanning tree (all PVST+, all rapid PVST+, or all MSTP).

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 IEEE 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 IEEE 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 IEEE 802.1Q VLAN of the trunk with the spanning-tree instance of the non-Cisco IEEE 802.1Q switch.

However, all PVST+ or rapid-PVST+ information is maintained by Cisco switches separated by a cloud of non-Cisco IEEE 802.1Q switches. The non-Cisco IEEE 802.1Q cloud separating the Cisco switches is treated as a single trunk link between the switches.

PVST+ is automatically enabled on IEEE 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 IEEE 802.1Q trunks, see [Chapter 1, “Configuring VLANs.”](#)

Spanning Tree and Switch Stacks

These statements are true when the switch stack is operating in PVST+ or rapid-PVST+ mode:

- A switch stack appears as a single spanning-tree node to the rest of the network, and all stack members use the same bridge ID for a given spanning tree. The bridge ID is derived from the MAC address of the stack master.
- When a new switch joins the stack, it sets its bridge ID to the stack-master bridge ID. If the newly added switch has the lowest ID and if the root path cost is the same among all stack members, the newly added switch becomes the stack root.
- When a stack member leaves the stack, spanning-tree reconvergence occurs within the stack (and possibly outside the stack). The remaining stack member with the lowest stack port ID becomes the stack root.
- If the stack master fails or leaves the stack, the stack members elect a new stack master, and all stack members change their bridge IDs of the spanning trees to the new master bridge ID.
- If the switch stack is the spanning-tree root and the stack master fails or leaves the stack, the stack members elect a new stack master, and a spanning-tree reconvergence occurs.
- If a neighboring switch external to the switch stack fails or is powered down, normal spanning-tree processing occurs. Spanning-tree reconvergence might occur as a result of losing a switch in the active topology.
- If a new switch external to the switch stack is added to the network, normal spanning-tree processing occurs. Spanning-tree reconvergence might occur as a result of adding a switch in the network.

For more information about switch stacks, see [Chapter 1, “Managing Switch Stacks.”](#)

Configuring Spanning-Tree Features

- [Default Spanning-Tree Configuration, page 1-13](#)
- [Spanning-Tree Configuration Guidelines, page 1-13](#)
- [Changing the Spanning-Tree Mode., page 1-15 \(required\)](#)
- [Disabling Spanning Tree, page 1-16 \(optional\)](#)
- [Configuring the Root Switch, page 1-16 \(optional\)](#)
- [Configuring a Secondary Root Switch, page 1-18 \(optional\)](#)
- [Configuring Port Priority, page 1-18 \(optional\)](#)
- [Configuring Path Cost, page 1-20 \(optional\)](#)
- [Configuring the Switch Priority of a VLAN, page 1-21 \(optional\)](#)
- [Configuring Spanning-Tree Timers, page 1-22 \(optional\)](#)

Default Spanning-Tree Configuration

Table 1-3 shows the default spanning-tree configuration.

Table 1-3 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 1-10.
Spanning-tree mode	PVST+. (Rapid PVST+ and MSTP are disabled.)
Switch priority	32768.
Spanning-tree port priority (configurable on a per-interface basis)	128.
Spanning-tree port cost (configurable on a per-interface basis)	1000 Mb/s: 4. 100 Mb/s: 19. 10 Mb/s: 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 Mb/s: 4. 100 Mb/s: 19. 10 Mb/s: 100.
Spanning-tree timers	Hello time: 2 seconds. Forward-delay time: 15 seconds. Maximum-aging time: 20 seconds. Transmit hold count: 6 BPDUs

Spanning-Tree Configuration Guidelines

Each stack member runs its own spanning tree, and the entire stack appears as a single switch to the rest of the network.

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 each switch stack. 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 1, “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.) All stack members run the same version of spanning tree. For information about the different spanning-tree modes and how they interoperate, see the [“Spanning-Tree Interoperability and Backward Compatibility” section on page 1-11](#).

For configuration guidelines about UplinkFast, BackboneFast, and cross-stack UplinkFast, see the [“Optional Spanning-Tree Configuration Guidelines” section on page 1-12](#).

**Caution**

Loop guard works only on point-to-point links. We recommend that each end of the link has a directly connected device that is running STP.

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. All stack members run the same version of spanning-tree. Note Stacking is supported only on Catalyst 2960-S switches running the LAN base image. <ul style="list-style-type: none"> • Select pvst to enable PVST+ (the default setting). • Select mst to enable MSTP (and RSTP). For more configuration steps, see Chapter 1, “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 6.
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 changes 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 IEEE 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+.
Step 7	show spanning-tree summary and show spanning-tree interface interface-id	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 1-10. 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 1-1 on page 1-5](#).)



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. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. 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.
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 a 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. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. 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 1-16 .
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.

**Note**

If your switch is a member of a switch stack, you must use the **spanning-tree [vlan *vlan-id*] cost *cost*** interface configuration command instead of the **spanning-tree [vlan *vlan-id*] port-priority *priority*** interface configuration command to select an interface to put in the forwarding state. Assign lower cost values to interfaces that you want selected first and higher cost values that you want selected last. For more information, see the “[Configuring Path Cost](#)” section on page 1-20.

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. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. 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 1-19.

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 interface number in the forwarding state and blocks the other interfaces.

Beginning in privileged EXEC mode, follow these steps to configure the cost of an interface. This procedure is optional.

	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. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. 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 1-19.

Configuring the Switch Priority of a VLAN

You can configure the switch priority and make it more likely that a standalone switch or a switch in the stack 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. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. 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 1-4 describes the timers that affect the entire spanning-tree performance.

Table 1-4 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.
Transmit hold count	Controls the number of BPDUs that can be sent before pausing for 1 second.

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. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. 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. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. 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. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. 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.

Configuring the Transmit Hold-Count

You can configure the BPDU burst size by changing the transmit hold count value.



Note

Changing this parameter to a higher value can have a significant impact on CPU utilization, especially in Rapid-PVST mode. Lowering this value can slow down convergence in certain scenarios. We recommend that you maintain the default setting.

Beginning in privileged EXEC mode, follow these steps to configure the transmit hold-count. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree transmit hold-count <i>value</i>	Configure the number of BPDUs that can be sent before pausing for 1 second. For <i>value</i> , the range is 1 to 20; the default is 6.
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 transmit hold-count** *value* 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 1-5](#):

Table 1-5 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, see the command reference for this release.



CHAPTER 1

Configuring MSTP

This chapter describes how to configure the Cisco implementation of the IEEE 802.1s Multiple STP (MSTP) on the Catalyst 2960 and 2960-S switches.



Note

The multiple spanning-tree (MST) implementation is based on the IEEE 802.1s standard. The MST implementations in Cisco IOS releases earlier than Cisco IOS Release 12.2(25)SED are prestandard.

The MSTP enables multiple VLANs to be mapped to the same spanning-tree instance, 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 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) IEEE 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 1, “Configuring STP.”](#) For information about other spanning-tree features such as Port Fast, UplinkFast, root guard, and so forth, see [Chapter 1, “Configuring Optional Spanning-Tree Features.”](#)

A switch stack appears as a single spanning-tree node to the rest of the network, and all stack members use the same switch ID. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.



Note

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release.

- [Understanding MSTP, page 1-2](#)
- [Understanding RSTP, page 1-9](#)

- [Configuring MSTP Features, page 1-14](#)
- [Displaying the MST Configuration and Status, page 1-27](#)

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.

- [Multiple Spanning-Tree Regions, page 1-2](#)
- [IST, CIST, and CST, page 1-3](#)
- [Hop Count, page 1-5](#)
- [Boundary Ports, page 1-6](#)
- [IEEE 802.1s Implementation, page 1-6](#)
- [“MSTP and Switch Stacks” section on page 1-8](#)
- [Interoperability with IEEE 802.1D STP, page 1-9](#)

For configuration information, see the [“Configuring MSTP Features” section on page 1-14](#).

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 1-1 on page 1-4](#).

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 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 65 spanning-tree instances. Instances can be identified by any number in the range from 0 to 4094. 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 4094.

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 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 by the spanning-tree algorithm running among switches that support the IEEE 802.1w, IEEE 802.1s, and IEEE 802.1D standards. 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 1-3 and the “Operations Between MST Regions” section on page 1-4.



Note

The implementation of the IEEE 802.1s standard, changes some of the terminology associated with MST implementations. For a summary of these changes, see [Table 1-1 on page 1-5](#).

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 CIST regional root (called the *IST master* before the implementation of the IEEE 802.1s standard) as shown in [Figure 1-1 on page 1-4](#). It is the switch within the region with the lowest switch ID and path cost to the CIST root. The CIST regional root is also the CIST root if there is only one region in the network. If the CIST root is outside the region, one of the MSTP switches at the boundary of the region is selected as the CIST regional root.

When an MSTP switch initializes, it sends BPDUs claiming itself as the root of the CIST and the CIST regional root, with both of the path costs to the CIST root and to the CIST regional root 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 switch ID, lower path cost, and so forth) than currently stored for the port, it relinquishes its claim as the CIST regional root.

During initialization, a region might have many subregions, each with its own CIST regional root. As switches receive superior IST information, they leave their old subregions and join the new subregion that contains the true CIST regional root. Thus all subregions shrink, except for the one that contains the true CIST regional root.

For correct operation, all switches in the MST region must agree on the same CIST regional root. Therefore, any two switches in the region only synchronize their port roles for an MST instance if they converge to a common CIST regional root.


Operations Between MST Regions

If there are multiple regions or legacy IEEE 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 CIST that encompasses the entire switched domain. The root of the subtree is the CIST regional root. The MST region appears as a virtual switch to adjacent STP switches and MST regions.

Figure 1-1 shows a network with three MST regions and a legacy IEEE 802.1D switch (D). The CIST regional root for region 1 (A) is also the CIST root. The CIST regional root for region 2 (B) and the CIST regional root for region 3 (C) are the roots for their respective subtrees within the CIST. The RSTP runs in all regions.

Figure 1-1 MST Regions, CIST Masters, and CST Root



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,

92983

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, and port VLAN priority) can be configured on both the CST instance and the MST instance.

MSTP switches use Version 3 RSTP BPDUs or IEEE 802.1D STP BPDUs to communicate with legacy IEEE 802.1D switches. MSTP switches use MSTP BPDUs to communicate with MSTP switches.

IEEE 802.1s Terminology

Some MST naming conventions used in Cisco's prestandard implementation have been changed to identify some *internal* or *regional* parameters. These parameters are significant only within an MST region, as opposed to external parameters that are relevant to the whole network. Because the CIST is the only spanning-tree instance that spans the whole network, only the CIST parameters require the external rather than the internal or regional qualifiers.

- The CIST root is the root switch for the unique instance that spans the whole network, the CIST.
- The CIST external root path cost is the cost to the CIST root. This cost is left unchanged within an MST region. Remember that an MST region looks like a single switch for the CIST. The CIST external root path cost is the root path cost calculated between these virtual switches and switches that do not belong to any region.
- The CIST regional root was called the IST master in the prestandard implementation. If the CIST root is in the region, the CIST regional root is the CIST root. Otherwise, the CIST regional root is the closest switch to the CIST root in the region. The CIST regional root acts as a root switch for the IST.
- The CIST internal root path cost is the cost to the CIST regional root in a region. This cost is only relevant to the IST, instance 0.

Table 1-1 on page 1-5 compares the IEEE standard and the Cisco prestandard terminology.

Table 1-1 Prestandard and Standard Terminology

IEEE Standard	Cisco Prestandard	Cisco Standard
CIST regional root	IST master	CIST regional root
CIST internal root path cost	IST master path cost	CIST internal path cost
CIST external root path cost	Root path cost	Root path cost
MSTI regional root	Instance root	Instance root
MSTI internal root path cost	Root path cost	Root path cost

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 (triggers 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 designated ports at the boundary.

Boundary Ports

In the Cisco prestandard implementation, a boundary port 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.

There is no definition of a boundary port in the IEEE 802.1s standard. The IEEE 802.1Q-2002 standard identifies two kinds of messages that a port can receive: internal (coming from the same region) and external. When a message is external, it is received only by the CIST. If the CIST role is root or alternate, or if the external BPDU is a topology change, it could have an impact on the MST instances. When a message is internal, the CIST part is received by the CIST, and each MST instance receives its respective M-record. The Cisco prestandard implementation treats a port that receives an external message as a boundary port. This means a port cannot receive a mix of internal and external messages.

An MST region includes both switches and LANs. A segment belongs to the region of its designated port. Therefore, a port in a different region than the designated port for a segment is a boundary port. This definition allows two ports internal to a region to share a segment with a port belonging to a different region, creating the possibility of receiving both internal and external messages on a port.

The primary change from the Cisco prestandard implementation is that a designated port is not defined as boundary, unless it is running in an STP-compatible mode.

**Note**

If there is a legacy STP switch on the segment, messages are always considered external.

The other change from the prestandard implementation is that the CIST regional root switch ID field is now inserted where an RSTP or legacy IEEE 802.1Q switch has the sender switch ID. The whole region performs like a single virtual switch by sending a consistent sender switch ID to neighboring switches. In this example, switch C would receive a BPDU with the same consistent sender switch ID of root, whether or not A or B is designated for the segment.

IEEE 802.1s Implementation

The Cisco implementation of the IEEE MST standard includes features required to meet the standard, as well as some of the desirable prestandard functionality that is not yet incorporated into the published standard.

Port Role Naming Change

The boundary role is no longer in the final MST standard, but this boundary concept is maintained in Cisco's implementation. However, an MST instance port at a boundary of the region might not follow the state of the corresponding CIST port. Two cases exist now:


- The boundary port is the root port of the CIST regional root—When the CIST instance port is proposed and is in sync, it can send back an agreement and move to the forwarding state only after all the corresponding MSTI ports are in sync (and thus forwarding). The MSTI ports now have a special *master* role.
- The boundary port is not the root port of the CIST regional root—The MSTI ports follow the state and role of the CIST port. The standard provides less information, and it might be difficult to understand why an MSTI port can be alternately blocking when it receives no BPDUs (MRecords). In this case, although the boundary role no longer exists, the **show** commands identify a port as boundary in the *type* column of the output.

Interoperation Between Legacy and Standard Switches

Because automatic detection of prestandard switches can fail, you can use an interface configuration command to identify prestandard ports. A region cannot be formed between a standard and a prestandard switch, but they can interoperate by using the CIST. Only the capability of load balancing over different instances is lost in that particular case. The CLI displays different flags depending on the port configuration when a port receives prestandard BPDUs. A syslog message also appears the first time a switch receives a prestandard BPDU on a port that has not been configured for prestandard BPDU transmission.

Figure 1-2 illustrates this scenario. Assume that A is a standard switch and B a prestandard switch, both configured to be in the same region. A is the root switch for the CIST, and thus B has a root port (BX) on segment X and an alternate port (BY) on segment Y. If segment Y flaps, and the port on BY becomes the alternate before sending out a single prestandard BPDU, AY cannot detect that a prestandard switch is connected to Y and continues to send standard BPDUs. The port BY is thus fixed in a boundary, and no load balancing is possible between A and B. The same problem exists on segment X, but B might transmit topology changes.

Figure 1-2 Standard and Prestandard Switch Interoperation



Note

We recommend that you minimize the interaction between standard and prestandard MST implementations.


Detecting Unidirectional Link Failure

This feature is not yet present in the IEEE MST standard, but it is included in this Cisco IOS release. The software checks the consistency of the port role and state in the received BPDUs to detect unidirectional link failures that could cause bridging loops.

When a designated port detects a conflict, it keeps its role, but reverts to discarding state because disrupting connectivity in case of inconsistency is preferable to opening a bridging loop.

Figure 1-3 illustrates a unidirectional link failure that typically creates a bridging loop. Switch A is the root switch, and its BPDUs are lost on the link leading to switch B. RSTP and MST BPDUs include the role and state of the sending port. With this information, switch A can detect that switch B does not react to the superior BPDUs it sends and that switch B is the designated, not root switch. As a result, switch A blocks (or keeps blocking) its port, thus preventing the bridging loop.

Figure 1-3 Detecting Unidirectional Link Failure



92722

MSTP and Switch Stacks



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

A switch stack appears as a single spanning-tree node to the rest of the network, and all stack members use the same switch ID for a given spanning tree. The switch ID is derived from the MAC address of the stack master.

If a switch that does not support MSTP is added to a switch stack that does support MSTP or the reverse, the switch is put into a version mismatch state. If possible, the switch is automatically upgraded or downgraded to the same version of software that is running on the switch stack.

When a new switch joins the stack, it sets its switch ID to the stack master switch ID. If the newly added switch has the lowest ID and if the root path cost is the same among all stack members, the newly added switch becomes the stack root. A topology change occurs if the newly added switch contains a better root port for the switch stack or a better designated port for the LAN connected to the stack. The newly added switch causes a topology change in the network if another switch connected to the newly added switch changes its root port or designated ports.

When a stack member leaves the stack, spanning-tree reconvergence occurs within the stack (and possibly outside the stack). The remaining stack member with the lowest stack port ID becomes the stack root.

If the stack master fails or leaves the stack, the stack members elect a new stack master, and all stack members change their switch IDs of the spanning trees to the new master switch ID.

For more information about switch stacks, see [Chapter 1, “Managing Switch Stacks.”](#)

Interoperability with IEEE 802.1D STP

A switch running MSTP supports a built-in protocol migration mechanism that enables it to interoperate with legacy IEEE 802.1D switches. If this switch receives a legacy IEEE 802.1D configuration BPDU (a BPDU with the protocol version set to 0), it sends only IEEE 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 IEEE 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 might also 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 IEEE 802.1D spanning tree).

- [Port Roles and the Active Topology, page 1-9](#)
- [Rapid Convergence, page 1-10](#)
- [Synchronization of Port Roles, page 1-11](#)
- [Bridge Protocol Data Unit Format and Processing, page 1-12](#)

For configuration information, see the “[Configuring MSTP Features](#)” section on page 1-14.

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 1-3. 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 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 IEEE 802.1D). The port state controls the operation of the forwarding and learning processes. [Table 1-2](#) provides a comparison of IEEE 802.1D and RSTP port states.

Table 1-2 Port State Comparison

Operational Status	STP Port State (IEEE 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 defines 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 1-4](#), 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.

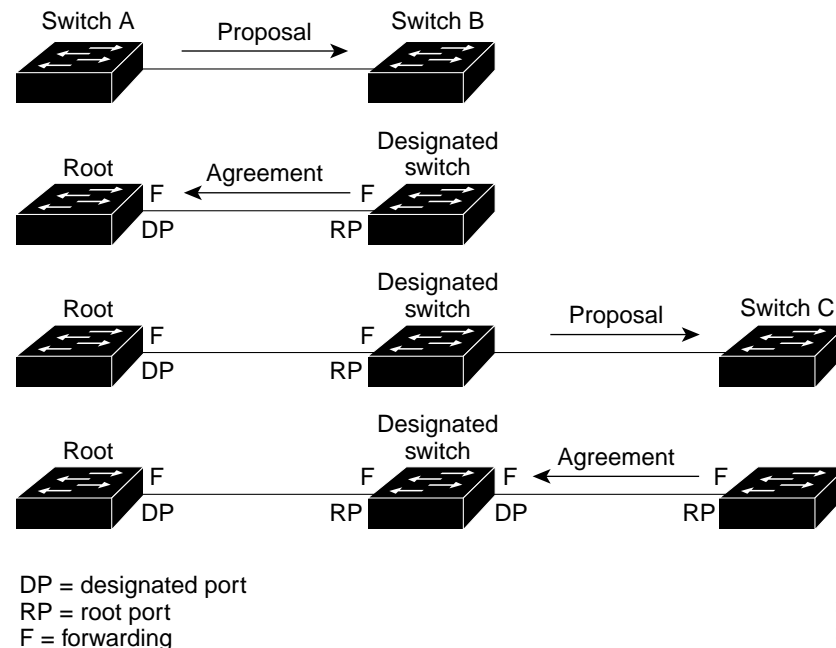
After receiving Switch B's agreement message, Switch A also immediately transitions its designated port to the forwarding state. No loops in the network are formed because Switch B blocked all of its nonedge ports and because there is a point-to-point link between Switches A and B.

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.

In a switch stack, the cross-stack rapid transition (CSRT) feature ensures that a stack member receives acknowledgments from all stack members during the proposal-agreement handshaking before moving the port to the forwarding state. CSRT is automatically enabled when the switch is in MST mode.

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 1-4 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 that 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 1-5](#).

Figure 1-5 Sequence of Events During Rapid Convergence



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 1-byte Version 1 Length field is set to zero, which means that no version 1 protocol information is present. [Table 1-3](#) shows the RSTP flag fields.

Table 1-3 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 IEEE 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 switch 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 IEEE 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 switch 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 IEEE 802.1D in handling spanning-tree topology changes.

- **Detection**—Unlike IEEE 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 deletes the learned information on all of its nonedge ports except on those from which it received the TC notification.
- **Notification**—Unlike IEEE 802.1D, which uses TCN BPDUs, the RSTP does not use them. However, for IEEE 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 IEEE 802.1D switch, it replies with an IEEE 802.1D configuration BPDU with the TCA bit set. However, if the TC-while timer (the same as the topology-change timer in IEEE 802.1D) is active on a root port connected to an IEEE 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 IEEE 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 IEEE 802.1D switches, RSTP selectively sends IEEE 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 IEEE 802.1D BPDU after the port migration-delay timer has expired, it assumes that it is connected to an IEEE 802.1D switch and starts using only IEEE 802.1D BPDUs. However, if the RSTP switch is using IEEE 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

- [Default MSTP Configuration, page 1-14](#)
- [MSTP Configuration Guidelines, page 1-15](#)
- [Specifying the MST Region Configuration and Enabling MSTP, page 1-16](#) (required)
- [Configuring the Root Switch, page 1-18](#) (optional)
- [Configuring a Secondary Root Switch, page 1-19](#) (optional)
- [Configuring Port Priority, page 1-20](#) (optional)
- [Configuring Path Cost, page 1-21](#) (optional)
- [Configuring the Switch Priority, page 1-22](#) (optional)
- [Configuring the Hello Time, page 1-23](#) (optional)
- [Configuring the Forwarding-Delay Time, page 1-24](#) (optional)
- [Configuring the Maximum-Aging Time, page 1-24](#) (optional)
- [Configuring the Maximum-Hop Count, page 1-25](#) (optional)
- [Specifying the Link Type to Ensure Rapid Transitions, page 1-25](#) (optional)
- [Designating the Neighbor Type, page 1-26](#) (optional)
- [Restarting the Protocol Migration Process, page 1-26](#) (optional)

Default MSTP Configuration

Table 1-4 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.

Table 1-4 Default MSTP Configuration (continued)

Feature	Default Setting
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 1-10.

MSTP Configuration Guidelines



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

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 stacked 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 stack supports up to 65 MST instances. The number of VLANs that can be mapped to a particular MST instance is unlimited.
- 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 1-11. For information on the recommended trunk port configuration, see the [“Interaction with Other Features”](#) section on page 1-15.
- All stack members run the same version of spanning tree (all PVST+, rapid PVST+, or MSTP). For more information, see the [“Spanning-Tree Interoperability and Backward Compatibility”](#) section on page 1-11.
- 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. You can achieve load balancing across a switch stack by manually configuring the path cost.
- 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.
- For configuration guidelines about UplinkFast, BackboneFast, and cross-stack UplinkFast, see the “Optional Spanning-Tree Configuration Guidelines” section on page 1-12.
- When the switch is in MST mode, it uses the long path-cost calculation method (32 bits) to compute the path cost values. With the long path-cost calculation method, these path cost values are supported:

Speed	Path Cost Value
10 Mb/s	2,000,000
100 Mb/s	200,000
1 Gb/s	20,000
10 Gb/s	2,000
100 Gb/s	200


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 only support up to 65 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.

	Command	Purpose
Step 3	instance <i>instance-id</i> vlan <i>vlan-range</i>	<p>Map VLANs to an MST instance.</p> <ul style="list-style-type: none"> For <i>instance-id</i>, the range is 0 to 4094. For vlan <i>vlan-range</i>, the range is 1 to 4094. <p>When you map VLANs to an MST instance, the mapping is incremental, and the VLANs specified in the command are added to or removed from the VLANs that were previously mapped.</p> <p>To specify a VLAN range, use a hyphen; for example, instance 1 vlan 1-63 maps VLANs 1 through 63 to MST instance 1.</p> <p>To specify a VLAN series, use a comma; for example, instance 1 vlan 10, 20, 30 maps VLANs 10, 20, and 30 to MST instance 1.</p>
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.
Step 8	spanning-tree mode mst	<p>Enable MSTP. RSTP is also enabled.</p> <p> Caution Changing spanning-tree modes can disrupt traffic because all spanning-tree instances are stopped for the previous mode and restarted in the new mode.</p> <p>You cannot run both MSTP and PVST+ or both MSTP and rapid PVST+ at the same time.</p>
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 switch 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 switch 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 1-1 on page 1-5](#).)

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.

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 [<i>diameter net-diameter</i> [<i>hello-time seconds</i>]]	Configure a switch as the root switch. <ul style="list-style-type: none"> For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 4094. (Optional) For diameter net-diameter, 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 seconds, 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 a 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 <i>instance-id</i> root secondary [diameter <i>net-diameter</i> [hello-time <i>seconds</i>]]	Configure a switch as the secondary root switch. <ul style="list-style-type: none"> For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 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. 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 1-18.
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 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.



Note

If your switch is a member of a switch stack, you must use the **spanning-tree mst [*instance-id*] cost *cost*** interface configuration command instead of the **spanning-tree mst [*instance-id*] port-priority *priority*** interface configuration command to select a port to put in the forwarding state. Assign lower cost values to ports that you want selected first and higher cost values to ports that you want selected last. For more information, see the [“Configuring Path Cost”](#) section on page 1-21.

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

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 6.
Step 3	spanning-tree mst <i>instance-id</i> port-priority <i>priority</i>	Configure the port priority. <ul style="list-style-type: none"> For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 4094. For <i>priority</i>, the range is 0 to 240 in increments of 16. The default is 128. The lower the number, the higher the priority. 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 6.
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. <ul style="list-style-type: none"> For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 4094. 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 ports that are in a link-up operative state. Otherwise, you can use the **show running-config** privileged EXEC command to confirm the configuration.

To return the interface to its default setting, use the **no spanning-tree mst** *instance-id* **cost** interface configuration command.

Configuring the Switch Priority

You can configure the switch priority and make it more likely that a standalone switch or a switch in the stack will be chosen as the root switch.

**Note**

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

**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 <i>instance-id</i> priority <i>priority</i>	Configure the switch priority. <ul style="list-style-type: none"> For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 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. 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 <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* 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.

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 255; 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 1-10](#).

By default, the link type is controlled from the duplex mode of the interface: a full-duplex port is considered to have a point-to-point connection; a half-duplex port is considered to have a shared connection. If you have a half-duplex link physically connected point-to-point to a single port on a remote switch running MSTP, you can override the default setting of the link type and enable rapid transitions to the forwarding state.

Beginning in privileged EXEC mode, follow these steps to override the default link-type setting. This procedure is optional.

	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, VLANs, and port-channel logical interfaces. The VLAN ID range is 1 to 4094. The port-channel range is 1 to 6.
Step 3	spanning-tree link-type point-to-point	Specify that the link type of a port is point-to-point.
Step 4	end	Return to privileged EXEC mode.
Step 5	show spanning-tree mst 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 the port to its default setting, use the **no spanning-tree link-type** interface configuration command.

Designating the Neighbor Type

A topology could contain both prestandard and IEEE 802.1s standard compliant devices. By default, ports can automatically detect prestandard devices, but they can still receive both standard and prestandard BPDUs. When there is a mismatch between a device and its neighbor, only the CIST runs on the interface.

You can choose to set a port to send only prestandard BPDUs. The prestandard flag appears in all the show commands, even if the port is in STP compatibility mode.

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 2	interface <i>interface-id</i>	Specify an interface to configure, and enter interface configuration mode. Valid interfaces include physical ports.
Step 3	spanning-tree mst pre-standard	Specify that the port can send only prestandard BPDUs.
Step 4	end	Return to privileged EXEC mode.
Step 5	show spanning-tree mst 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 the port to its default setting, use the **no spanning-tree mst prestandard** 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 IEEE 802.1D switches. If this switch receives a legacy IEEE 802.1D configuration BPDU (a BPDU with the protocol version set to 0), it sends only IEEE 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 IEEE 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 interface, use the **clear spanning-tree detected-protocols interface** *interface-id* privileged EXEC command.

Displaying the MST Configuration and Status

To display the spanning-tree status, use one or more of the privileged EXEC commands in [Table 1-5](#):

Table 1-5 *Commands for Displaying MST Status*

Command	Purpose
show spanning-tree mst configuration	Displays the MST region configuration.
show spanning-tree mst configuration digest	Displays the MD5 digest included in the current MSTCI.
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, see the command reference for this release.



CHAPTER 1

Configuring Optional Spanning-Tree Features

This chapter describes how to configure optional spanning-tree features on the Catalyst 2960 and 2960-S switches. 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 stack is running the Multiple Spanning Tree Protocol (MSTP) or the rapid per-VLAN spanning-tree plus (rapid-PVST+) protocol. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

For information on configuring the PVST+ and rapid PVST+, see [Chapter 1, “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 1, “Configuring MSTP.”](#)



Note

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release.

- [Understanding Optional Spanning-Tree Features, page 1-1](#)
- [Configuring Optional Spanning-Tree Features, page 1-11](#)
- [Displaying the Spanning-Tree Status, page 1-19](#)

Understanding Optional Spanning-Tree Features

- [Understanding Port Fast, page 1-2](#)
- [Understanding BPDU Guard, page 1-2](#)
- [Understanding BPDU Filtering, page 1-3](#)
- [Understanding UplinkFast, page 1-3](#)
- [Understanding BackboneFast, page 1-7](#)
- [Understanding EtherChannel Guard, page 1-10](#)
- [Understanding Root Guard, page 1-10](#)
- [Understanding Loop Guard, page 1-11](#)

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 1-1](#), 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.

You can enable this feature by using the **spanning-tree portfast** interface configuration or the **spanning-tree portfast default** global configuration command.

Figure 1-1 Port Fast-Enabled Interfaces



Understanding BPDU Guard

The BPDU guard feature can be globally enabled on the switch or can be enabled per port, but the feature operates with some differences.

At the global level, you enable BPDU guard on Port Fast-enabled ports by using the **spanning-tree portfast bpduguard default** global configuration command. Spanning tree shuts down ports that are in a Port Fast-operational state if any BPDU is received on them. In a valid configuration, Port Fast-enabled ports do not receive BPDUs. Receiving a BPDU on a Port Fast-enabled port means an invalid configuration, such as the connection of an unauthorized device, and the BPDU guard feature puts the port in the error-disabled state. When this happens, the switch shuts down the entire port on which the violation occurred.

To prevent the port from shutting down, you can use the **errdisable detect cause bpduguard shutdown vlan** global configuration command to shut down just the offending VLAN on the port where the violation occurred.

At the interface level, you enable BPDU guard on any port by using the **spanning-tree bpduguard enable** interface configuration command without also enabling the Port Fast feature. When the port receives a BPDU, it is put in the error-disabled state.

The BPDU guard feature provides a secure response to invalid configurations because you must manually put the 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.

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.

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 1-2](#) 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 1-2 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.

When the spanning tree reconfigures the new root port, other interfaces flood the network with multicast packets, one for each address that was learned on the interface. You can limit these bursts of multicast traffic by reducing the max-update-rate parameter (the default for this parameter is 150 packets per second). However, if you enter zero, station-learning frames are not generated, so the spanning-tree topology converges more slowly after a loss of connectivity.


**Note**

UplinkFast is most useful in wiring-closet switches at the access or edge of the network. It is not appropriate for backbone devices. This feature might not be useful for other types of applications.

UplinkFast provides fast convergence after a direct link failure and achieves load balancing between redundant Layer 2 links using uplink groups. An uplink group is a set of Layer 2 interfaces (per VLAN), only one of which is forwarding at any given time. Specifically, an uplink group consists of the root port (which is forwarding) and a set of blocked ports, except for self-looping ports. The uplink group provides an alternate path in case the currently forwarding link fails.


Figure 1-3 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 1-3 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 1-4. This change takes approximately 1 to 5 seconds.

Figure 1-4 UplinkFast Example After Direct Link Failure



Understanding Cross-Stack UplinkFast



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

For Catalyst 37502975 switches, the UplinkFast feature is the cross-stack UplinkFast feature. Cross-stack UplinkFast (CSUF) provides a fast spanning-tree transition (fast convergence in less than 1 second under normal network conditions) across a switch stack. During the fast transition, an alternate redundant link on the switch stack is placed in the forwarding state without causing temporary spanning-tree loops or loss of connectivity to the backbone. With this feature, you can have a redundant and resilient network in some configurations. CSUF is automatically enabled when you enable the UplinkFast feature by using the **spanning-tree uplinkfast** global configuration command.

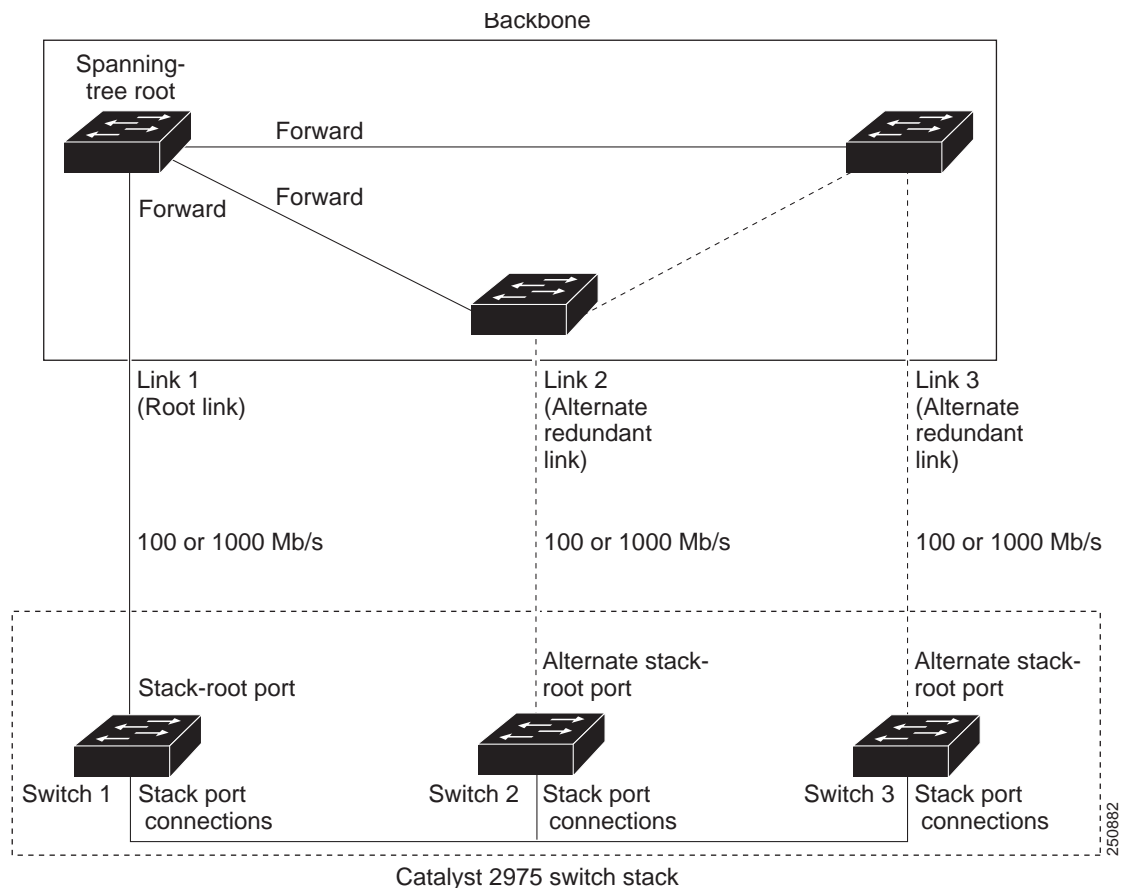
CSUF might not provide a fast transition all the time; in these cases, the normal spanning-tree transition occurs, completing in 30 to 40 seconds. For more information, see the “[Events that Cause Fast Convergence](#)” section on page 1-7.

How CSUF Works

CSUF ensures that one link in the stack is elected as the path to the root. As shown in [Figure 1-5](#), the stack-root port on Switch 1 provides the path to the root of the spanning tree. The alternate stack-root ports on Switches 2 and 3 can provide an alternate path to the spanning-tree root if the current stack-root switch fails or if its link to the spanning-tree root fails.

Link 1, the root link, is in the spanning-tree forwarding state. Links 2 and 3 are alternate redundant links that are in the spanning-tree blocking state. If Switch 1 fails, if its stack-root port fails, or if Link 1 fails, CSUF selects either the alternate stack-root port on Switch 2 or Switch 3 and puts it into the forwarding state in less than 1 second.

Figure 1-5 Cross-Stack UplinkFast Topology



When certain link loss or spanning-tree events occur (described in [“Events that Cause Fast Convergence” section on page 1-7](#)), the Fast Uplink Transition Protocol uses the neighbor list to send fast-transition requests to stack members.

The switch sending the fast-transition request needs to do a fast transition to the forwarding state of a port that it has chosen as the root port, and it must obtain an acknowledgement from each stack switch before performing the fast transition.

Each switch in the stack decides if the sending switch is a better choice than itself to be the stack root of this spanning-tree instance by comparing the root, cost, and bridge ID. If the sending switch is the best choice as the stack root, each switch in the stack returns an acknowledgement; otherwise, it sends a fast-transition request. The sending switch then has not received acknowledgements from all stack switches.

When acknowledgements are received from all stack switches, the Fast Uplink Transition Protocol on the sending switch immediately transitions its alternate stack-root port to the forwarding state. If acknowledgements from all stack switches are not obtained by the sending switch, the normal spanning-tree transitions (blocking, listening, learning, and forwarding) take place, and the spanning-tree topology converges at its normal rate ($2 * \text{forward-delay time} + \text{max-age time}$).

The Fast Uplink Transition Protocol is implemented on a per-VLAN basis and affects only one spanning-tree instance at a time.

Events that Cause Fast Convergence

Depending on the network event or failure, the CSUF fast convergence might or might not occur.

Fast convergence (less than 1 second under normal network conditions) occurs under these circumstances:

- The stack-root port link fails.

If two switches in the stack have alternate paths to the root, only one of the switches performs the fast transition.

- The failed link, which connects the stack root to the spanning-tree root, recovers.
- A network reconfiguration causes a new stack-root switch to be selected.
- A network reconfiguration causes a new port on the current stack-root switch to be chosen as the stack-root port.



Note

The fast transition might not occur if multiple events occur simultaneously. For example, if a stack member is powered off, and at the same time, the link connecting the stack root to the spanning-tree root comes back up, the normal spanning-tree convergence occurs.

Normal spanning-tree convergence (30 to 40 seconds) occurs under these conditions:

- The stack-root switch is powered off, or the software failed.
- The stack-root switch, which was powered off or failed, is powered on.
- A new switch, which might become the stack root, is added to the stack.

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.

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 find 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 to learn if any stack member has an alternate root to the root switch and waits for an RLQ reply from other switches in the network and in the stack. The switch sends the RLQ request on all alternate paths and waits for an RLQ reply from other switches in the network.

When a stack member receives an RLQ reply from a nonstack member on a blocked interface and the reply is destined for another nonstacked switch, it forwards the reply packet, regardless of the spanning-tree interface state.

When a stack member receives an RLQ reply from a nonstack member and the response is destined for the stack, the stack member forwards the reply so that all the other stack members receive it.


**Note**

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

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.

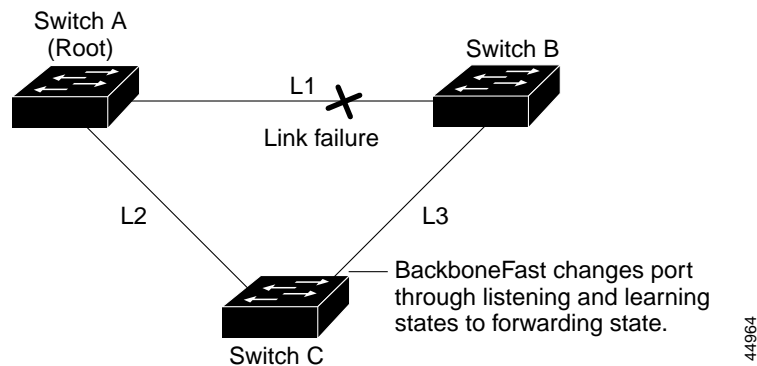
Figure 1-6 shows an example topology with no link failures. Switch A, the root switch, connects directly to Switch B over link L1 and to Switch C over link L2. The Layer 2 interface on Switch C that connects directly to Switch B is in the blocking state.

Figure 1-6 BackboneFast Example Before Indirect Link Failure




If link L1 fails as shown in [Figure 1-7](#), 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. The root-switch election takes approximately 30 seconds, twice the Forward Delay time if the default Forward Delay time of 15 seconds is set. [Figure 1-7](#) shows how BackboneFast reconfigures the topology to account for the failure of link L1.

Figure 1-7 BackboneFast Example After Indirect Link Failure



If a new switch is introduced into a shared-medium topology as shown in [Figure 1-8](#), 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 1-8 Adding a Switch in a Shared-Medium Topology



Understanding EtherChannel Guard

You can use EtherChannel guard to detect an EtherChannel misconfiguration between the switch and a connected device. A misconfiguration can occur if the switch interfaces are configured in an EtherChannel, but the interfaces on the other device are not. A misconfiguration can also occur if the channel parameters are not the same at both ends of the EtherChannel. For EtherChannel configuration guidelines, see the [“EtherChannel Configuration Guidelines”](#) section on page 1-11.

If the switch detects a misconfiguration on the other device, EtherChannel guard places the switch interfaces in the error-disabled state, and displays an error message.

You can enable this feature by using the **spanning-tree etherchannel guard misconfig** global configuration command.

Understanding Root Guard

The Layer 2 network of a service provider (SP) can include many connections to switches that are not owned by the SP. In such a topology, the spanning tree can reconfigure itself and select a *customer switch* as the root switch, as shown in [Figure 1-9](#). 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 IEEE 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.

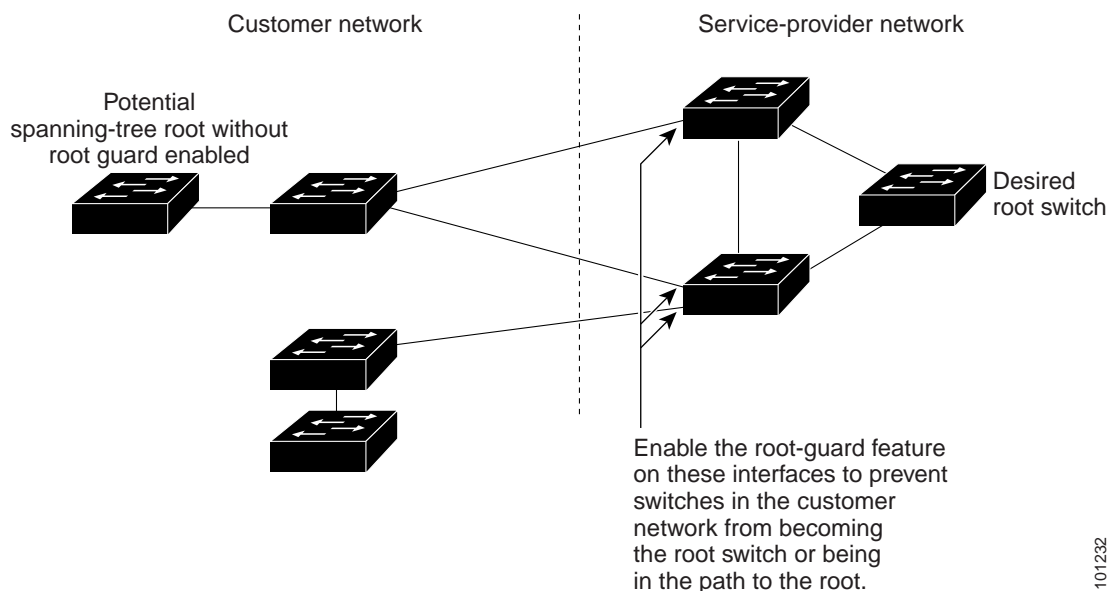
You can enable this feature by using the **spanning-tree guard root** interface configuration command.



Caution

Misuse of the root-guard feature can cause a loss of connectivity.

Figure 1-9 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.

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

- [Default Optional Spanning-Tree Configuration, page 1-12](#)
- [Optional Spanning-Tree Configuration Guidelines, page 1-12](#)
- [Enabling Port Fast, page 1-12 \(optional\)](#)

- [Enabling BPDU Guard, page 1-14](#) (optional)
- [Enabling BPDU Filtering, page 1-15](#) (optional)
- [Enabling UplinkFast for Use with Redundant Links, page 1-16](#) (optional)
- [Enabling Cross-Stack UplinkFast, page 1-16](#) (optional)
- [Enabling BackboneFast, page 1-17](#) (optional)
- [Enabling EtherChannel Guard, page 1-17](#) (optional)
- [Enabling Root Guard, page 1-18](#) (optional)
- [Enabling Loop Guard, page 1-18](#) (optional)

Default Optional Spanning-Tree Configuration

Table 1-1 shows the default optional spanning-tree configuration.

Table 1-1 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. (On Catalyst 2960-S switches, the UplinkFast feature is the CSUF feature.)
BackboneFast	Globally disabled.
EtherChannel guard	Globally enabled.
Root guard	Disabled on all interfaces.
Loop guard	Disabled on all interfaces.

Optional Spanning-Tree Configuration Guidelines

You can configure PortFast, BPDU guard, BPDU filtering, EtherChannel guard, root guard, or loop guard if your switch is running PVST+, rapid PVST+, or MSTP.

You can configure the UplinkFast, the BackboneFast, or the cross-stack UplinkFast feature for rapid PVST+ or for the MSTP, but the feature remains disabled (inactive) until you change the spanning-tree mode to PVST+.

You can configure the UplinkFast or the BackboneFast feature for rapid PVST+ or for the MSTP, but the feature remains disabled (inactive) until you change the spanning-tree mode to PVST+.

Enabling Port Fast

An interface with the Port Fast feature enabled is moved directly to the spanning-tree forwarding state without waiting for the standard forward-time delay.


**Caution**

Use Port Fast *only* when connecting a single end station to an access or trunk port. Enabling this feature on 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 1, “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]	<p>Enable Port Fast on an access port connected to a single workstation or server. By specifying the trunk keyword, you can enable Port Fast on a trunk port.</p> <p>Note To enable Port Fast on trunk ports, you must use the spanning-tree portfast trunk interface configuration command. The spanning-tree portfast command will not work on trunk ports.</p> <p> Caution Make sure that there are no loops in the network between the trunk port and the workstation or server before you enable Port Fast on a trunk port.</p> <p>By default, Port Fast is disabled on all interfaces.</p>
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 ports that are Port Fast-enabled (the ports are in a Port Fast-operational state), spanning tree continues to run on the ports. They remain up unless they receive a BPDU.

In a valid configuration, Port Fast-enabled ports do not receive BPDUs. Receiving a BPDU on a Port Fast-enabled port means an invalid configuration, such as the connection of an unauthorized device, and the BPDU guard feature puts the port in the error-disabled state. When this happens, the switch shuts down the entire port on which the violation occurred.

To prevent the port from shutting down, you can use the **errdisable detect cause bpduguard shutdown vlan** global configuration command to shut down just the offending VLAN on the port where the violation occurred.

The BPDU guard feature provides a secure response to invalid configurations because you must manually put the port back in service. Use the BPDU guard feature in a service-provider network to prevent an access port from participating in the spanning tree.



Caution

Configure Port Fast only on ports that connect to end stations; otherwise, an accidental topology loop could cause a data packet loop and disrupt switch and network operation.

You also can use the **spanning-tree bpduguard enable** interface configuration command to enable BPDU guard on any port without also enabling the Port Fast feature. When the port receives a BPDU, it is put in the error-disabled state.

You can enable the BPDU guard feature if your switch is running PVST+, rapid PVST+, or MSTP.

Beginning in privileged EXEC mode, follow these steps to globally enable the BPDU guard feature. This procedure is optional.

	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 stack. You cannot configure UplinkFast on an individual VLAN.

You can configure the UplinkFast or the CSUF feature for rapid PVST+ or for the MSTP, but the feature remains disabled (inactive) until you change the spanning-tree mode to PVST+.

Beginning in privileged EXEC mode, follow these steps to enable UplinkFast and CSUF. 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. When you enter this command, CSUF also is enabled on all nonstack port interfaces.
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 Cross-Stack UplinkFast

When you enable or disable the UplinkFast feature by using the **spanning-tree uplinkfast** global configuration command, CSUF is automatically globally enabled or disabled on nonstack port interfaces.

For more information, see the “[Enabling UplinkFast for Use with Redundant Links](#)” section on page 1-16.

To disable UplinkFast on the switch and all its VLANs, use the **no spanning-tree uplinkfast** global configuration 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.

You can configure the BackboneFast feature for rapid PVST+ or for the MSTP, but the feature remains disabled (inactive) until you change the spanning-tree mode to PVST+.

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 EtherChannel Guard

You can enable EtherChannel guard to detect an EtherChannel misconfiguration if your switch is running PVST+, rapid PVST+, or MSTP.

Beginning in privileged EXEC mode, follow these steps to enable EtherChannel guard. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree etherchannel guard misconfig	Enable EtherChannel guard.
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 EtherChannel guard feature, use the **no spanning-tree etherchannel guard misconfig** global configuration command.

You can use the **show interfaces status err-disabled** privileged EXEC command to show which switch ports are disabled because of an EtherChannel misconfiguration. On the remote device, you can enter the **show etherchannel summary** privileged EXEC command to verify the EtherChannel configuration.

After the configuration is corrected, enter the **shutdown** and **no shutdown** interface configuration commands on the port-channel interfaces that were misconfigured.

Enabling Root Guard

Root guard enabled on an interface applies to all the VLANs to which the interface belongs. Do not enable the root guard on interfaces to be used by the UplinkFast feature. With UplinkFast, the backup interfaces (in the blocked state) replace the root port in the case of a failure. However, if root guard is also enabled, all the backup interfaces used by the UplinkFast feature are placed in the root-inconsistent state (blocked) and are prevented from reaching the forwarding state.



Note

You cannot enable both root guard and loop guard at the same time.

You can enable this feature if your switch is running PVST+, rapid PVST+, or MSTP.

Beginning in privileged EXEC mode, follow these steps to enable root guard on an interface. This procedure is optional.

	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.
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 or show spanning-tree mst	Verify which interfaces are alternate or root ports.
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 1-2](#):

Table 1-2 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, see the command reference for this release.

■ Displaying the Spanning-Tree Status



CHAPTER 1

Configuring Flex Links and the MAC Address-Table Move Update Feature



Note

To use Flex Links and the MAC address-table move update feature, the switch must be running the LAN Base image.

This chapter describes how to configure Flex Links, a pair of interfaces on the Catalyst 2960 and 2960-S switches that provide a mutual backup. It also describes how to configure the MAC address-table move update feature, also referred to as the Flex Links bidirectional fast convergence feature. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release.

- [Understanding Flex Links and the MAC Address-Table Move Update, page 1-1](#)
- [Configuring Flex Links and the MAC Address-Table Move Update, page 1-7](#)
- [Monitoring Flex Links and the MAC Address-Table Move Update, page 1-14](#)

Understanding Flex Links and the MAC Address-Table Move Update

- [Flex Links, page 1-2](#)
- [VLAN Flex Link Load Balancing and Support, page 1-2](#)
- [Flex Link Multicast Fast Convergence, page 1-3](#)
- [MAC Address-Table Move Update, page 1-6](#)

Flex Links


Flex Links are a pair of a Layer 2 interfaces (switch ports or port channels) where one interface is configured to act as a backup to the other. The feature provides an alternative solution to the Spanning Tree Protocol (STP). Users can disable STP and still retain basic link redundancy. Flex Links are typically configured in service provider or enterprise networks where customers do not want to run STP on the switch. If the switch is running STP, Flex Links is not necessary because STP already provides link-level redundancy or backup.

You configure Flex Links on one Layer 2 interface (the active link) by assigning another Layer 2 interface as the Flex Link or backup link. The Flex Link can be on the same switch or on another switch in the stack. When one of the links is up and forwarding traffic, the other link is in standby mode, ready to begin forwarding traffic if the other link shuts down. At any given time, only one of the interfaces is in the linkup state and forwarding traffic. If the primary link shuts down, the standby link starts forwarding traffic. When the active link comes back up, it goes into standby mode and does not forward traffic. STP is disabled on Flex Link interfaces.

In [Figure 1-1](#), ports 1 and 2 on switch A are connected to uplink switches B and C. Because they are configured as Flex Links, only one of the interfaces is forwarding traffic; the other is in standby mode. If port 1 is the active link, it begins forwarding traffic between port 1 and switch B; the link between port 2 (the backup link) and switch C is not forwarding traffic. If port 1 goes down, port 2 comes up and starts forwarding traffic to switch C. When port 1 comes back up, it goes into standby mode and does not forward traffic; port 2 continues forwarding traffic.

You can also choose to configure a preemption mechanism, specifying the preferred port for forwarding traffic. For example, in the example in [Figure 1-1](#), you can configure the Flex Links pair with preemption mode. In the scenario shown, when port 1 comes back up and has more bandwidth than port 2, port 1 begins forwarding traffic after 60 seconds. Port 2 becomes the standby port. You do this by entering the interface configuration **switchport backup interface preemption mode bandwidth** and **switchport backup interface preemption delay** commands.

Figure 1-1 Flex Links Configuration Example



If a primary (forwarding) link goes down, a trap notifies the network management stations. If the standby link goes down, a trap notifies the users.


Flex Links are supported only on Layer 2 ports and port channels, not on VLANs.

VLAN Flex Link Load Balancing and Support

VLAN Flex Link load-balancing allows you to configure a Flex Link pair so that both ports simultaneously forward the traffic for some mutually exclusive VLANs. For example, if Flex Link ports are configured for 1-100 VLANs, the traffic of the first 50 VLANs can be forwarded on one port and the rest on the other port. If one of the ports fail, the other active port forwards all the traffic. When the failed

port comes back up, it resumes forwarding traffic in the preferred VLANs. This way, apart from providing the redundancy, this Flex Link pair can be used for load balancing. Also, Flex Link VLAN load-balancing does not impose any restrictions on uplink switches.

Figure 1-2 VLAN Flex Links Load Balancing Configuration Example



201398

Flex Link Multicast Fast Convergence



Note

To use Flex Link Multicast Fast Convergence, the switch must be running the LAN Base image.

Flex Link Multicast Fast Convergence reduces the multicast traffic convergence time after a Flex Link failure. This is implemented by a combination of these solutions:

- [Learning the Other Flex Link Port as the mrouter Port, page 1-3](#)
- [Generating IGMP Reports, page 1-4](#)
- [Leaking IGMP Reports, page 1-4](#)
- [Configuration Examples, page 1-4](#)

Learning the Other Flex Link Port as the mrouter Port

In a typical multicast network, there is a querier for each VLAN. A switch deployed at the edge of a network has one of its Flex Link ports receiving queries. Flex Link ports are also always forwarding at any given time.

A port that receives queries is added as an *mrouter* port on the switch. An mrouter port is part of all the multicast groups learned by the switch. After a changeover, queries are received by the other Flex Link port. The other Flex Link port is then learned as the mrouter port. After changeover, multicast traffic then flows through the other Flex Link port. To achieve faster convergence of traffic, both Flex Link ports are learned as mrouter ports whenever either Flex Link port is learned as the mrouter port. Both Flex Link ports are always part of multicast groups.

Though both Flex Link ports are part of the groups in normal operation mode, all traffic on the backup port is blocked. So the normal multicast data flow is not affected by the addition of the backup port as an mrouter port. When the changeover happens, the backup port is unblocked, allowing the traffic to flow. In this case, the upstream multicast data flows as soon as the backup port is unblocked.

Generating IGMP Reports

When the backup link comes up after the changeover, the upstream new distribution switch does not start forwarding multicast data, because the port on the upstream router, which is connected to the blocked Flex Link port, is not part of any multicast group. The reports for the multicast groups were not forwarded by the downstream switch because the backup link is blocked. The data does not flow on this port, until it learns the multicast groups, which occurs only after it receives reports.

The reports are sent by hosts when a general query is received, and a general query is sent within 60 seconds in normal scenarios. When the backup link starts forwarding, to achieve faster convergence of multicast data, the downstream switch immediately sends proxy reports for all the learned groups on this port without waiting for a general query.

Leaking IGMP Reports

To achieve multicast traffic convergence with minimal loss, a redundant data path must be set up before the Flex Link active link goes down. This can be achieved by leaking only IGMP report packets on the Flex Link backup link. These leaked IGMP report messages are processed by upstream distribution routers, so multicast data traffic gets forwarded to the backup interface. Because all incoming traffic on the backup interface is dropped at the ingress of the access switch, no duplicate multicast traffic is received by the host. When the Flex Link active link fails, the access switch starts accepting traffic from the backup link immediately. The only disadvantage of this scheme is that it consumes bandwidth on the link between the distribution switches and on the backup link between the distribution and access switches. This feature is disabled by default and can be configured by using the **switchport backup interface *interface-id* multicast fast-convergence** command.

When this feature has been enabled at changeover, the switch does not generate the proxy reports on the backup port, which became the forwarding port.

Configuration Examples

These are configuration examples for learning the other Flex Link port as the mrouter port when Flex Link is configured on Gigabit Ethernet0/11 and Gigabit Ethernet0/12, with output for the **show interfaces switchport backup** command:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabithernet1/0/11
Switch(config)# interface gigabithernet0/11
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport mode trunk
Switch(config-if)# switchport backup interface gigabithernet0/12
Switch(config-if)# exit
Switch(config)# interface gigabithernet1/0/12
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport mode trunk
Switch(config-if)# end
Switch# show interfaces switchport backup detail
Switch Backup Interface Pairs:
Active Interface Backup Interface State
GigabitEthernet0/11 GigabitEthernet0/12 Active Up/Backup Standby
Preemption Mode : off
Multicast Fast Convergence : Off
Bandwidth : 100000 Kbit (Gi0/11), 100000 Kbit (Gi0/12)
Mac Address Move Update Vlan : auto
```


This output shows a querier for VLANs 1 and 401, with their queries reaching the switch through Gigabit Ethernet0/11:

```
Switch# show ip igmp snooping querier
Vlan    IP Address    IGMP Version    Port
-----
1       1.1.1.1       v2               Gi0/11
401     41.41.41.1    v2               Gi0/11
```

Here is output for the **show ip igmp snooping mrouter** command for VLANs 1 and 401:

```
Switch# show ip igmp snooping mrouter
Vlan    ports
-----
1       Gi1/0/11(dynamic), Gi1/0/12(dynamic)
401     Gi1/0/11(dynamic), Gi1/0/12(dynamic)
```

Similarly, both Flex Link ports are part of learned groups. In this example, Gigabit Ethernet0/11 is a receiver/host in VLAN 1, which is interested in two multicast groups:

```
Switch# show ip igmp snooping groups
Vlan    Group        Type    Version    Port List
-----
1       228.1.5.1    igmp    v2         Gi1/0/11, Gi1/0/12, Gi2/0/11
1       228.1.5.2    igmp    v2         Gi1/0/11, Gi1/0/12, Gi2/0/11
```

When a host responds to the general query, the switch forwards this report on all the mrouter ports. In this example, when a host sends a report for the group 228.1.5.1, it is forwarded only on Gigabit Ethernet0/11, because the backup port Gigabit Ethernet0/12 is blocked. When the active link, Gigabit Ethernet0/11, goes down, the backup port, Gigabit Ethernet0/12, begins forwarding.

As soon as this port starts forwarding, the switch sends proxy reports for the groups 228.1.5.1 and 228.1.5.2 on behalf of the host. The upstream router learns the groups and starts forwarding multicast data. This is the default behavior of Flex Link. This behavior changes when the user configures fast convergence using the **switchport backup interface gigabitEthernet 0/12 multicast fast-convergence** command. This example shows turning on this feature:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitEthernet 0/11
Switch(config-if)# switchport backup interface gigabitEthernet 0/12 multicast
fast-convergence
Switch(config-if)# exit
Switch# show interfaces switchport backup detail
Switch Backup Interface Pairs:
Active          Interface          Backup Interface State
-----
GigabitEthernet0/11 GigabitEthernet0/12 Active Up/Backup Standby
Preemption Mode : off
Multicast Fast Convergence : On
Bandwidth : 100000 Kbit (Gi0/11), 100000 Kbit (Gi0/12)
Mac Address Move Update Vlan : auto
```

This output shows a querier for VLAN 1 and 401 with their queries reaching the switch through Gigabit Ethernet0/11:

```
Switch# show ip igmp snooping querier
Vlan    IP Address    IGMP Version    Port
-----
1       1.1.1.1       v2               Gi0/11
401     41.41.41.1    v2               Gi0/11
```

This is output for the **show ip igmp snooping mrouter** command for VLAN 1 and 401:

```
Switch# show ip igmp snooping mrouter
Vlan      ports
----      -
1         Gi0/11(dynamic), Gi0/12(dynamic)
401      Gi10/11(dynamic), Gi0/12(dynamic)
```

Similarly, both the Flex Link ports are a part of the learned groups. In this example, Gigabit Ethernet0/11 is a receiver/host in VLAN 1, which is interested in two multicast groups:

```
Switch# show ip igmp snooping groups
Vlan  Group      Type   Version  Port List
-----
1     228.1.5.1   igmp   v2       Gi1/0/11, Gi1/0/12, Gi2/0/11
1     228.1.5.2   igmp   v2       Gi1/0/11, Gi1/0/12, Gi2/0/11
```

Whenever a host responds to the general query, the switch forwards this report on all the mrouter ports. When you turn on this feature through the command-line port, and when a report is forwarded by the switch on GigabitEthernet0/11, it is also leaked to the backup port GigabitEthernet0/12. The upstream router learns the groups and starts forwarding multicast data, which is dropped at the ingress because GigabitEthernet0/12 is blocked. When the active link, GigabitEthernet0/11, goes down, the backup port, GigabitEthernet0/12, begins forwarding. You do not need to send any proxy reports because the multicast data is already being forwarded by the upstream router. By leaking reports to the backup port, a redundant multicast path has been set up, and the time taken for the multicast traffic convergence is minimal.

MAC Address-Table Move Update

The MAC address-table move update feature allows the switch to provide rapid bidirectional convergence when a primary (forwarding) link goes down and the standby link begins forwarding traffic.

In [Figure 1-3](#), switch A is an access switch, and ports 1 and 2 on switch A are connected to uplink switches B and D through a Flex Link pair. Port 1 is forwarding traffic, and port 2 is in the backup state. Traffic from the PC to the server is forwarded from port 1 to port 3. The MAC address of the PC has been learned on port 3 of switch C. Traffic from the server to the PC is forwarded from port 3 to port 1.


If the MAC address-table move update feature is not configured and port 1 goes down, port 2 starts forwarding traffic. However, for a short time, switch C keeps forwarding traffic from the server to the PC through port 3, and the PC does not get the traffic because port 1 is down. If switch C removes the MAC address of the PC on port 3 and relearns it on port 4, traffic can then be forwarded from the server to the PC through port 2.

If the MAC address-table move update feature is configured and enabled on the switches in [Figure 1-3](#) and port 1 goes down, port 2 starts forwarding traffic from the PC to the server. The switch sends a MAC address-table move update packet from port 2. Switch C gets this packet on port 4 and immediately learns the MAC address of the PC on port 4, which reduces the reconvergence time.

You can configure the access switch, switch A, to *send* MAC address-table move update messages. You can also configure the uplink switches B, C, and D to *get* and process the MAC address-table move update messages. When switch C gets a MAC address-table move update message from switch A, switch C learns the MAC address of the PC on port 4. Switch C updates the MAC address table, including the forwarding table entry for the PC.

Switch A does not need to wait for the MAC address-table update. The switch detects a failure on port 1 and immediately starts forwarding server traffic from port 2, the new forwarding port. This change occurs in 100 milliseconds (ms). The PC is directly connected to switch A, and the connection status does not change. Switch A does not need to update the PC entry in the MAC address table.

Figure 1-3 MAC Address-Table Move Update Example



Configuring Flex Links and the MAC Address-Table Move Update

- [Default Configuration, page 1-7](#)
- [Configuration Guidelines, page 1-8](#)
- [Configuring Flex Links, page 1-8](#)
- [Configuring VLAN Load Balancing on Flex Links, page 1-10](#)
- [Configuring the MAC Address-Table Move Update Feature, page 1-12](#)

Default Configuration

The Flex Links are not configured, and there are no backup interfaces defined.

The preemption mode is off.

The preemption delay is 35 seconds.

The MAC address-table move update feature is not configured on the switch.

Configuration Guidelines

Follow these guidelines to configure Flex Links:

- You can configure up to 16 backup links.
- You can configure only one Flex Link backup link for any active link, and it must be a different interface from the active interface.
- An interface can belong to only one Flex Link pair. An interface can be a backup link for only one active link. An active link cannot belong to another Flex Link pair.
- Neither of the links can be a port that belongs to an EtherChannel. However, you can configure two port channels (EtherChannel logical interfaces) as Flex Links, and you can configure a port channel and a physical interface as Flex Links, with either the port channel or the physical interface as the active link.
- A backup link does not have to be the same type (Fast Ethernet, Gigabit Ethernet, or port channel) as the active link. However, you should configure both Flex Links with similar characteristics so that there are no loops or changes in behavior if the standby link begins to forward traffic.
- STP is disabled on Flex Link ports. A Flex Link port does not participate in STP, even if the VLANs present on the port are configured for STP. When STP is not enabled, be sure that there are no loops in the configured topology. Once the Flex Link configurations are removed, STP is re-enabled on the ports.

Follow these guidelines to configure VLAN load balancing on the Flex Links feature:

- For Flex Link VLAN load balancing, you must choose the preferred VLANs on the backup interface.
- You cannot configure a preemption mechanism and VLAN load balancing for the same Flex Links pair.

Follow these guidelines to configure the MAC address-table move update feature:

- You can enable and configure this feature on the access switch to *send* the MAC address-table move updates.
- You can enable and configure this feature on the uplink switches to *receive* the MAC address-table move updates.

Configuring Flex Links

Beginning in privileged EXEC mode, follow these steps to configure a pair of Flex Links:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface, and enter interface configuration mode. The interface can be a physical Layer 2 interface or a port channel (logical interface). The port-channel range is 1 to 6.

	Command	Purpose
Step 3	switchport backup interface <i>interface-id</i>	Configure a physical Layer 2 interface (or port channel) as part of a Flex Link pair with the interface. When one link is forwarding traffic, the other interface is in standby mode.
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces [<i>interface-id</i>] switchport backup	Verify the configuration.
Step 6	copy running-config startup config	(Optional) Save your entries in the switch startup configuration file.

To disable a Flex Link backup interface, use the **no switchport backup interface** *interface-id* interface configuration command.

This example shows how to configure an interface with a backup interface and to verify the configuration:

```
Switch# configure terminal
Switch(conf)# interface gigabitethernet0/1
Switch(conf-if)# switchport backup interface gigabitethernet0/2
Switch(conf-if)# end

Switch# show interfaces switchport backup
Switch Backup Interface Pairs:

Active Interface      Backup Interface      State
-----
GigabitEthernet0/1   GigabitEthernet0/3   Active Standby/Backup Up
Vlans Preferred on Active Interface: 1-3,5-4094
Vlans Preferred on Backup Interface: 4
```

Beginning in privileged EXEC mode, follow these steps to configure a preemption scheme for a pair of Flex Links:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface, and enter interface configuration mode. The interface can be a physical Layer 2 interface or a port channel (logical interface). The port-channel range is 1 to 6.
Step 3	switchport backup interface <i>interface-id</i>	Configure a physical Layer 2 interface (or port channel) as part of a Flex Links pair with the interface. When one link is forwarding traffic, the other interface is in standby mode.
Step 4	switchport backup interface <i>interface-id</i> preemption mode [forced bandwidth off]	Configure a preemption mechanism and delay for a Flex Link interface pair. You can configure the preemption as: <ul style="list-style-type: none"> Forced—the active interface always preempts the backup. Bandwidth—the interface with the higher bandwidth always acts as the active interface. Off—no preemption happens from active to backup.

	Command	Purpose
Step 5	switchport backup interface <i>interface-id</i> preempt mode delay <i>delay-time</i>	Configure the time delay until a port preempts another port. Note Setting a delay time only works with forced and bandwidth modes.
Step 6	end	Return to privileged EXEC mode.
Step 7	show interfaces [<i>interface-id</i>] switchport backup	Verify the configuration.
Step 8	copy running-config startup config	(Optional) Save your entries in the switch startup configuration file.

To remove a preemption scheme, use the **no switchport backup interface *interface-id* preempt mode** interface configuration command. To reset the delay time to the default, use the **no switchport backup interface *interface-id* preempt delay** interface configuration command.

This example shows how to configure the preemption mode as *forced* for a backup interface pair and to verify the configuration:

Catalyst 2960-S switch:

```
Switch# configure terminal
```

Catalyst 2960 switch:

```
Switch(conf)# interface gigabitethernet0/1
Switch(conf-if)#switchport backup interface gigabitethernet0/2 preempt mode forced
Switch(conf-if)#switchport backup interface gigabitethernet0/2 preempt delay 50
Switch(conf-if)# end
```

```
Switch# show interfaces switchport backup detail
Active Interface Backup Interface State
```

Catalyst 2960 switch:

```
GigabitEthernet0/21 GigabitEthernet0/2 Active Up/Backup Standby
Interface Pair : Gi0/1, Gi0/2
Preemption Mode : forced
Preemption Delay : 50 seconds
Bandwidth : 100000 Kbit (Gi0/1), 100000 Kbit (Gi0/2)
Mac Address Move Update Vlan : auto
```

Configuring VLAN Load Balancing on Flex Links

Beginning in privileged EXEC mode, follow these steps to configure VLAN load balancing on Flex Links:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface, and enter interface configuration mode. The interface can be a physical Layer 2 interface or a port channel (logical interface). The port-channel range is 1 to 6.

	Command	Purpose
Step 3	switchport backup interface <i>interface-id</i> prefer vlan <i>vlan-range</i>	Configure a physical Layer 2 interface (or port channel) as part of a Flex Links pair with the interface, and specify the VLANs carried on the interface. The VLAN ID range is 1 to 4094.
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces [<i>interface-id</i>] switchport backup	Verify the configuration.
Step 6	copy running-config startup config	(Optional) Save your entries in the switch startup configuration file.

To disable the VLAN load balancing feature, use the **no switchport backup interface** *interface-id* **prefer vlan** *vlan-range* interface configuration command.

In the following example, VLANs 1 to 50, 60, and 100 to 120 are configured on the switch:

Catalyst 2960-S switch:

Catalyst 2960 switch:

```
Switch(config)#interface gigabitethernet 0/6
Switch(config-if)#switchport backup interface gigabitethernet 0/8 prefer vlan 60,100-120
```

When both interfaces are up, Gi2/0/8 forwards traffic for VLANs 60 and 100 to 120, and Gi0/6 forwards traffic for VLANs 1 to 50.

```
Switch#show interfaces switchport backup
Switch Backup Interface Pairs:
```

```
Active Interface      Backup Interface      State
-----
GigabitEthernet0/6   GigabitEthernet0/8   Active Up/Backup Up
Vlans Preferred on Active Interface: 1-50
Vlans Preferred on Backup Interface: 60, 100-120
```

When a Flex Link interface goes down (LINK_DOWN), VLANs preferred on this interface are moved to the peer interface of the Flex Link pair. In this example, if interface Gi2/0/6 goes down, Gi2/0/8 carries all VLANs of the Flex Link pair.

```
Switch# show interfaces switchport backup
Switch Backup Interface Pairs:
```

```
Active Interface      Backup Interface      State
-----
GigabitEthernet0/6   GigabitEthernet0/8   Active Down/Backup Up
Vlans Preferred on Active Interface: 1-50
Vlans Preferred on Backup Interface: 60, 100-120
```

When a Flex Link interface comes up, VLANs preferred on this interface are blocked on the peer interface and moved to the forwarding state on the interface that has just come up. In this example, if interface Gi2/0/6 comes up, VLANs preferred on this interface are blocked on the peer interface Gi0/8 and forwarded on Gi2/0/6.

```
Switch#show interfaces switchport backup
```

```
Switch Backup Interface Pairs:
```

```
Active Interface      Backup Interface      State
-----
GigabitEthernet0/6   GigabitEthernet0/8   Active Up/Backup Up
```

```
Vlans Preferred on Active Interface: 1-50
```

```
Vlans Preferred on Backup Interface: 60, 100-120
```

```
Switch#show interfaces switchport backup detail
```

```
Switch Backup Interface Pairs:
```

```
Active Interface      Backup Interface      State
-----
FastEthernet0/3       FastEthernet0/4       Active Down/Backup Up
```

```
Vlans Preferred on Active Interface: 1-2,5-4094
```

```
Vlans Preferred on Backup Interface: 3-4
```

```
Preemption Mode : off
```

```
Bandwidth : 10000 Kbit (Fa0/3), 100000 Kbit (Fa0/4)
```

```
Mac Address Move Update Vlan : auto
```

Configuring the MAC Address-Table Move Update Feature

This section contains this information:

- Configuring a switch to send MAC address-table move updates
- Configuring a switch to get MAC address-table move updates

Beginning in privileged EXEC mode, follow these steps to configure an access switch to send MAC address-table move updates:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface, and enter interface configuration mode. The interface can be a physical Layer 2 interface or a port channel (logical interface). The port-channel range is 1 to 6.
Step 3	switchport backup interface <i>interface-id</i> or switchport backup interface <i>interface-id</i> mmu primary vlan <i>vlan-id</i>	Configure a physical Layer 2 interface (or port channel), as part of a Flex Link pair with the interface. The MAC address-table move update VLAN is the lowest VLAN ID on the interface. Configure a physical Layer 2 interface (or port channel) and specify the VLAN ID on the interface, which is used for sending the MAC address-table move update. When one link is forwarding traffic, the other interface is in standby mode.

	Command	Purpose
Step 4	end	Return to global configuration mode.
Step 5	mac address-table move update transmit	Enable the access switch to send MAC address-table move updates to other switches in the network if the primary link goes down and the switch starts forwarding traffic through the standby link.
Step 6	end	Return to privileged EXEC mode.
Step 7	show mac address-table move update	Verify the configuration.
Step 8	copy running-config startup config	(Optional) Save your entries in the switch startup configuration file.

To disable the MAC address-table move update feature, use the **no mac address-table move update transmit** interface configuration command. To display the MAC address-table move update information, use the **show mac address-table move update** privileged EXEC command.

This example shows how to configure an access switch to send MAC address-table move update messages:

```
Switch(conf)# interface gigabitethernet0/1
Switch(conf-if)# switchport backup interface gigabitethernet0/2 mmu primary vlan 2
Switch(conf-if)# exit
Switch(conf)# mac address-table move update transmit
Switch(conf)# end
```

This example shows how to verify the configuration:

```
Switch# show mac-address-table move update
Switch-ID : 010b.4630.1780
Dst mac-address : 0180.c200.0010
Vlans/Macs supported : 1023/8320
Default/Current settings: Rcv Off/On, Xmt Off/On
Max packets per min : Rcv 40, Xmt 60
Rcv packet count : 5
Rcv conforming packet count : 5
Rcv invalid packet count : 0
Rcv packet count this min : 0
Rcv threshold exceed count : 0
Rcv last sequence# this min : 0
Rcv last interface : Po2
Rcv last src-mac-address : 000b.462d.c502
Rcv last switch-ID : 0403.fd6a.8700
Xmt packet count : 0
Xmt packet count this min : 0
Xmt threshold exceed count : 0
Xmt pak buf unavail cnt : 0
Xmt last interface : None
```

Beginning in privileged EXEC mode, follow these steps to configure a switch to get and process MAC address-table move update messages:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mac address-table move update receive	Enable the switch to get and process the MAC address-table move updates.
Step 3	end	Return to privileged EXEC mode.

	Command	Purpose
Step 4	show mac address-table move update	Verify the configuration.
Step 5	copy running-config startup config	(Optional) Save your entries in the switch startup configuration file.

To disable the MAC address-table move update feature, use the **no mac address-table move update receive** configuration command. To display the MAC address-table move update information, use the **show mac address-table move update** privileged EXEC command.

This example shows how to configure a switch to get and process MAC address-table move update messages:

```
Switch# configure terminal
Switch(conf)# mac address-table move update receive
Switch(conf)# end
```

Monitoring Flex Links and the MAC Address-Table Move Update

Table 1-1 shows the privileged EXEC commands for monitoring the Flex Links configuration and the MAC address-table move update information.

Table 1-1 Flex Links and MAC Address-Table Move Update Monitoring Commands

Command	Purpose
show interfaces [interface-id] switchport backup	Displays the Flex Link backup interface configured for an interface or all the configured Flex Links and the state of each active and backup interface (up or standby mode). When VLAN load balancing is enabled, the output displays the preferred VLANs on Active and Backup interfaces.
show mac address-table move update	Displays the MAC address-table move update information on the switch.



CHAPTER 1

Configuring DHCP and IP Source Guard Features

This chapter describes how to configure DHCP snooping and option-82 data insertion, and the DHCP server port-based address allocation features on the Catalyst 2960 and 2960-S switches. It also describes how to configure the IP source guard feature.



Note

To use the IP source guard feature, the switch must be running the LAN Base image. Stacking is supported only Catalyst 2960-S switches.



Note

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release, and see the “DHCP Commands” section in the *Cisco IOS IP Command Reference, Volume 1 of 3: Addressing and Services, Release 12.4* on Cisco.com.

- [Understanding DHCP Snooping, page 1-1](#)
- [Configuring DHCP Snooping, page 1-8](#)
- [Displaying DHCP Snooping Information, page 1-13](#)
- [Understanding IP Source Guard, page 1-13](#)
- [Configuring IP Source Guard, page 1-15](#)
- [Displaying IP Source Guard Information, page 1-20](#)
- [Understanding DHCP Server Port-Based Address Allocation, page 1-21](#)
- [Configuring DHCP Server Port-Based Address Allocation, page 1-21](#)
- [Displaying DHCP Server Port-Based Address Allocation, page 1-24](#)

Understanding DHCP Snooping

DHCP is widely used in LAN environments to dynamically assign host IP addresses from a centralized server, which significantly reduces the overhead of administration of IP addresses. DHCP also helps conserve the limited IP address space because IP addresses no longer need to be permanently assigned to hosts; only those hosts that are connected to the network consume IP addresses.

- [DHCP Server, page 1-2](#)
- [DHCP Relay Agent, page 1-2](#)

- [DHCP Snooping, page 1-2](#)
- [Option-82 Data Insertion, page 1-3](#)
- [DHCP Snooping Binding Database, page 1-6](#)
- [DHCP Snooping and Switch Stacks, page 1-7](#)

For information about the DHCP client, see the “*Configuring DHCP*” section of the “*IP Addressing and Services*” section of the *Cisco IOS IP Configuration Guide, Release 12.4* on Cisco.com.

DHCP Server

The DHCP server assigns IP addresses from specified address pools on a switch or router to DHCP clients and manages them. If the DHCP server cannot give the DHCP client the requested configuration parameters from its database, it forwards the request to one or more secondary DHCP servers defined by the network administrator.

DHCP Relay Agent

A DHCP relay agent is a Layer 3 device that forwards DHCP packets between clients and servers. Relay agents forward requests and replies between clients and servers when they are not on the same physical subnet. Relay agent forwarding is different from the normal Layer 2 forwarding, in which IP datagrams are switched transparently between networks. Relay agents receive DHCP messages and generate new DHCP messages to send on output interfaces.

DHCP Snooping

DHCP snooping is a DHCP security feature that provides network security by filtering untrusted DHCP messages and by building and maintaining a DHCP snooping binding database, also referred to as a DHCP snooping binding table.

DHCP snooping acts like a firewall between untrusted hosts and DHCP servers. You use DHCP snooping to differentiate between untrusted interfaces connected to the end user and trusted interfaces connected to the DHCP server or another switch.

**Note**

For DHCP snooping to function properly, all DHCP servers must be connected to the switch through trusted interfaces.

An untrusted DHCP message is a message that is received from outside the network or firewall. When you use DHCP snooping in a service-provider environment, an untrusted message is sent from a device that is not in the service-provider network, such as a customer’s switch. Messages from unknown devices are untrusted because they can be sources of traffic attacks.

The DHCP snooping binding database has the MAC address, the IP address, the lease time, the binding type, the VLAN number, and the interface information that corresponds to the local untrusted interfaces of a switch. It does not have information regarding hosts interconnected with a trusted interface.

In a service-provider network, a trusted interface is connected to a port on a device in the same network. An untrusted interface is connected to an untrusted interface in the network or to an interface on a device that is not in the network.

When a switch receives a packet on an untrusted interface and the interface belongs to a VLAN in which DHCP snooping is enabled, the switch compares the source MAC address and the DHCP client hardware address. If the addresses match (the default), the switch forwards the packet. If the addresses do not match, the switch drops the packet.

The switch drops a DHCP packet when one of these situations occurs:

- A packet from a DHCP server, such as a DHCP OFFER, DHCP ACK, DHCP NAK, or DHCP LEASE QUERY packet, is received from outside the network or firewall.
- A packet is received on an untrusted interface, and the source MAC address and the DHCP client hardware address do not match.
- The switch receives a DHCP RELEASE or DHCP DECLINE broadcast message that has a MAC address in the DHCP snooping binding database, but the interface information in the binding database does not match the interface on which the message was received.
- A DHCP relay agent forwards a DHCP packet that includes a relay-agent IP address that is not 0.0.0.0, or the relay agent forwards a packet that includes option-82 information to an untrusted port.

If the switch is an aggregation switch supporting DHCP snooping and is connected to an edge switch that is inserting DHCP option-82 information, the switch drops packets with option-82 information when packets are received on an untrusted interface. If DHCP snooping is enabled and packets are received on a trusted port, the aggregation switch does not learn the DHCP snooping bindings for connected devices and cannot build a complete DHCP snooping binding database.

When an aggregation switch can be connected to an edge switch through an untrusted interface and you enter the **ip dhcp snooping information option allow-untrusted** global configuration command, the aggregation switch accepts packets with option-82 information from the edge switch. The aggregation switch learns the bindings for hosts connected through an untrusted switch interface. The DHCP security features, such as dynamic ARP inspection or IP source guard, can still be enabled on the aggregation switch while the switch receives packets with option-82 information on untrusted input interfaces to which hosts are connected. The port on the edge switch that connects to the aggregation switch must be configured as a trusted interface.

Option-82 Data Insertion

In residential, metropolitan Ethernet-access environments, DHCP can centrally manage the IP address assignments for a large number of subscribers. When the DHCP option-82 feature is enabled on the switch, a subscriber device is identified by the switch port through which it connects to the network (in addition to its MAC address). Multiple hosts on the subscriber LAN can be connected to the same port on the access switch and are uniquely identified.




Note

The DHCP option-82 feature is supported only when DHCP snooping is globally enabled and on the VLANs to which subscriber devices using this feature are assigned.

Figure 1-1 is an example of a metropolitan Ethernet network in which a centralized DHCP server assigns IP addresses to subscribers connected to the switch at the access layer. Because the DHCP clients and their associated DHCP server do not reside on the same IP network or subnet, a DHCP relay agent (the Catalyst switch) is configured with a helper address to enable broadcast forwarding and to transfer DHCP messages between the clients and the server.

Figure 1-1 DHCP Relay Agent in a Metropolitan Ethernet Network



When you enable the DHCP snooping information option 82 on the switch, this sequence of events occurs:

- The host (DHCP client) generates a DHCP request and broadcasts it on the network.
- When the switch receives the DHCP request, it adds the option-82 information in the packet. The remote-ID suboption is the switch MAC address, and the circuit-ID suboption is the port identifier, **vlan-mod-port**, from which the packet is received.
- If the IP address of the relay agent is configured, the switch adds this IP address in the DHCP packet.
- The switch forwards the DHCP request that includes the option-82 field to the DHCP server.
- The DHCP server receives the packet. If the server is option-82-capable, it can use the remote ID, the circuit ID, or both to assign IP addresses and implement policies, such as restricting the number of IP addresses that can be assigned to a single remote ID or circuit ID. Then the DHCP server echoes the option-82 field in the DHCP reply.
- The DHCP server unicasts the reply to the switch if the request was relayed to the server by the switch. The switch verifies that it originally inserted the option-82 data by inspecting the remote ID and possibly the circuit ID fields. The switch removes the option-82 field and forwards the packet to the switch port that connects to the DHCP client that sent the DHCP request.

When the described sequence of events occurs, the values in these fields in [Figure 1-2](#) do not change:


- Circuit-ID suboption fields
 - Suboption type
 - Length of the suboption type
 - Circuit-ID type
 - Length of the circuit-ID type
- Remote-ID suboption fields
 - Suboption type
 - Length of the suboption type
 - Remote-ID type
 - Length of the remote-ID type

In the port field of the circuit-ID suboption, the port numbers start at 3. For example, on a switch with 24 10/100 ports and small form-factor pluggable (SFP) module slots, port 3 is the Fast Ethernet x/0/1 port, port 4 is the Fast Ethernet x/0/2 port, and so forth, where x is the stack member number. Port 27 is the SFP module slot x/0/1, and so forth.

Figure 1-2 shows the packet formats for the remote-ID suboption and the circuit-ID suboption. For the circuit-ID suboption, the module number corresponds to the switch number in the stack. The switch uses the packet formats when you globally enable DHCP snooping and enter the **ip dhcp snooping information option** global configuration command.

Figure 1-2 Suboption Packet Formats

Circuit ID Suboption Frame Format



Remote ID Suboption Frame Format





Figure 1-3 shows the packet formats for user-configured remote-ID and circuit-ID suboptions. The switch uses these packet formats when DHCP snooping is globally enabled and when the **ip dhcp snooping information option format remote-id** global configuration command and the **ip dhcp snooping vlan information option format-type circuit-id string** interface configuration command are entered.

The values for these fields in the packets change from the default values when you configure the remote-ID and circuit-ID suboptions:

- Circuit-ID suboption fields
 - The circuit-ID type is 1.
 - The length values are variable, depending on the length of the string that you configure.
- Remote-ID suboption fields
 - The remote-ID type is 1.
 - The length values are variable, depending on the length of the string that you configure.

Figure 1-3 User-Configured Suboption Packet Formats

Circuit ID Suboption Frame Format (for user-configured string):**Remote ID Suboption Frame Format (for user-configured string):**

DHCP Snooping Binding Database

When DHCP snooping is enabled, the switch uses the DHCP snooping binding database to store information about untrusted interfaces. The database can have up to 64,000 bindings.

Each database entry (*binding*) has an IP address, an associated MAC address, the lease time (in hexadecimal format), the interface to which the binding applies, and the VLAN to which the interface belongs. The database agent stores the bindings in a file at a configured location. At the end of each entry is a checksum that accounts for all the bytes from the start of the file through all the bytes associated with the entry. Each entry is 72 bytes, followed by a space and then the checksum value.

To keep the bindings when the switch reloads, you must use the DHCP snooping database agent. If the agent is disabled, dynamic ARP inspection or IP source guard is enabled, and the DHCP snooping binding database has dynamic bindings, the switch loses its connectivity. If the agent is disabled and only DHCP snooping is enabled, the switch does not lose its connectivity, but DHCP snooping might not prevent DHCP spoofing attacks.

When reloading, the switch reads the binding file to build the DHCP snooping binding database. The switch updates the file when the database changes.

When a switch learns of new bindings or when it loses bindings, the switch immediately updates the entries in the database. The switch also updates the entries in the binding file. The frequency at which the file is updated is based on a configurable delay, and the updates are batched. If the file is not updated in a specified time (set by the `write-delay` and `abort-timeout` values), the update stops.

This is the format of the file with bindings:

```
<initial-checksum>
TYPE DHCP-SNOOPING
VERSION 1
BEGIN
<entry-1> <checksum-1>
<entry-2> <checksum-1-2>
...
...
<entry-n> <checksum-1-2-...-n>
END
```

Each entry in the file is tagged with a checksum value that the switch uses to verify the entries when it reads the file. The *initial-checksum* entry on the first line distinguishes entries associated with the latest file update from entries associated with a previous file update.

This is an example of a binding file:

```
2bb4c2a1
TYPE DHCP-SNOOPING
VERSION 1
BEGIN
192.1.168.1 3 0003.47d8.c91f 2BB6488E interface-id 21ae5fbb
192.1.168.3 3 0003.44d6.c52f 2BB648EB interface-id 1bdb223f
192.1.168.2 3 0003.47d9.c8f1 2BB648AB interface-id 584a38f0
END
```

When the switch starts and the calculated checksum value equals the stored checksum value, the switch reads entries from the binding file and adds the bindings to its DHCP snooping binding database. The switch ignores an entry when one of these situations occurs:

- The switch reads the entry and the calculated checksum value does not equal the stored checksum value. The entry and the ones following it are ignored.
- An entry has an expired lease time (the switch might not remove a binding entry when the lease time expires).
- The interface in the entry no longer exists on the system.
- The interface is a routed interface or a DHCP snooping-trusted interface.

DHCP Snooping and Switch Stacks

DHCP snooping is managed on the stack master. When a new switch joins the stack, the switch receives the DHCP snooping configuration from the stack master. When a member leaves the stack, all DHCP snooping address bindings associated with the switch age out.

All snooping statistics are generated on the stack master. If a new stack master is elected, the statistics counters reset.

When a stack merge occurs, all DHCP snooping bindings in the stack master are lost if it is no longer the stack master. With a stack partition, the existing stack master is unchanged, and the bindings belonging to the partitioned switches age out. The new master of the partitioned stack begins processing the new incoming DHCP packets. For more information about switch stacks, see [Chapter 1, “Managing Switch Stacks.”](#)

Configuring DHCP Snooping

- [Default DHCP Snooping Configuration](#), page 1-8
- [DHCP Snooping Configuration Guidelines](#), page 1-8
- [Configuring the DHCP Relay Agent](#), page 1-10
- [Enabling DHCP Snooping and Option 82](#), page 1-10
- [Enabling the DHCP Snooping Binding Database Agent](#), page 1-12

Default DHCP Snooping Configuration

Table 1-1 shows the default DHCP snooping configuration.

Table 1-1 Default DHCP Snooping Configuration

Feature	Default Setting
DHCP server	Enabled in Cisco IOS software, requires configuration ¹
DHCP relay agent	Enabled ²
DHCP packet forwarding address	None configured
Checking the relay agent information	Enabled (invalid messages are dropped) ²
DHCP relay agent forwarding policy	Replace the existing relay agent information ²
DHCP snooping enabled globally	Disabled
DHCP snooping information option	Enabled
DHCP snooping option to accept packets on untrusted input interfaces ³	Disabled
DHCP snooping limit rate	None configured
DHCP snooping trust	Untrusted
DHCP snooping VLAN	Disabled
DHCP snooping MAC address verification	Enabled
DHCP snooping binding database agent	Enabled in Cisco IOS software, requires configuration. This feature is operational only when a destination is configured.

1. The switch responds to DHCP requests only if it is configured as a DHCP server.
2. The switch relays DHCP packets only if the IP address of the DHCP server is configured on the SVI of the DHCP client.
3. Use this feature when the switch is an aggregation switch that receives packets with option-82 information from an edge switch.

DHCP Snooping Configuration Guidelines

- You must globally enable DHCP snooping on the switch.
- DHCP snooping is not active until DHCP snooping is enabled on a VLAN.
- Before globally enabling DHCP snooping on the switch, make sure that the devices acting as the DHCP server and the DHCP relay agent are configured and enabled.

- Before configuring the DHCP snooping information option on your switch, be sure to configure the device that is acting as the DHCP server. For example, you must specify the IP addresses that the DHCP server can assign or exclude, or you must configure DHCP options for these devices.
- When configuring a large number of circuit IDs on a switch, consider the impact of lengthy character strings on the NVRAM or the flash memory. If the circuit-ID configurations, combined with other data, exceed the capacity of the NVRAM or the flash memory, an error message appears.
- Before configuring the DHCP relay agent on your switch, make sure to configure the device that is acting as the DHCP server. For example, you must specify the IP addresses that the DHCP server can assign or exclude, configure DHCP options for devices, or set up the DHCP database agent.
- If the DHCP relay agent is enabled but DHCP snooping is disabled, the DHCP option-82 data insertion feature is not supported.
- If a switch port is connected to a DHCP server, configure a port as trusted by entering the **ip dhcp snooping trust** interface configuration command.
- If a switch port is connected to a DHCP client, configure a port as untrusted by entering the **no ip dhcp snooping trust** interface configuration command.
- Follow these guidelines when configuring the DHCP snooping binding database:
 - Because both NVRAM and the flash memory have limited storage capacity, we recommend that you store the binding file on a TFTP server.
 - For network-based URLs (such as TFTP and FTP), you must create an empty file at the configured URL before the switch can write bindings to the binding file at that URL. See the documentation for your TFTP server to determine whether you must first create an empty file on the server; some TFTP servers cannot be configured this way.
 - To ensure that the lease time in the database is accurate, we recommend that you enable and configure NTP. For more information, see the [“Configuring Time and Date Manually” section on page 1-5](#).
 - If NTP is configured, the switch writes binding changes to the binding file only when the switch system clock is synchronized with NTP.
- Do not enter the **ip dhcp snooping information option allow-untrusted** command on an aggregation switch to which an untrusted device is connected. If you enter this command, an untrusted device might spoof the option-82 information.
- You can display DHCP snooping statistics by entering the **show ip dhcp snooping statistics** user EXEC command, and you can clear the snooping statistics counters by entering the **clear ip dhcp snooping statistics** privileged EXEC command.

**Note**

Do not enable Dynamic Host Configuration Protocol (DHCP) snooping on RSPAN VLANs. If DHCP snooping is enabled on RSPAN VLANs, DHCP packets might not reach the RSPAN destination port.

Configuring the DHCP Relay Agent

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

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	service dhcp	Enable the DHCP server and relay agent on your switch. By default, this feature is enabled.
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 the DHCP server and relay agent, use the **no service dhcp** global configuration command.

See the “*Configuring DHCP*” section of the “IP Addressing and Services” section of the *Cisco IOS IP Configuration Guide, Release 12.4* on Cisco.com for these procedures:

- Checking (validating) the relay agent information
- Configuring the relay agent forwarding policy

Enabling DHCP Snooping and Option 82

Beginning in privileged EXEC mode, follow these steps to enable DHCP snooping on the switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip dhcp snooping	Enable DHCP snooping globally.
Step 3	ip dhcp snooping vlan <i>vlan-range</i>	Enable DHCP snooping on a VLAN or range of VLANs. The range is 1 to 4094. You can enter a single VLAN ID identified by VLAN ID number, a series of VLAN IDs separated by commas, a range of VLAN IDs separated by hyphens, or a range of VLAN IDs separated by entering the starting and ending VLAN IDs separated by a space.
Step 4	ip dhcp snooping information option	Enable the switch to insert and to remove DHCP relay information (option-82 field) in forwarded DHCP request messages to the DHCP server. This is the default setting.
Step 5	ip dhcp snooping information option allow-untrusted	(Optional) If the switch is an aggregation switch connected to an edge switch, enable the switch to accept incoming DHCP snooping packets with option-82 information from the edge switch. The default setting is disabled. Note Enter this command only on aggregation switches that are connected to trusted devices.
Step 6	interface <i>interface-id</i>	Specify the interface to be configured, and enter interface configuration mode.

	Command	Purpose
Step 7	ip dhcp snooping trust	(Optional) Configure the interface as trusted or as untrusted. Use the no keyword to configure an interface to receive messages from an untrusted client. The default setting is untrusted.
Step 8	ip dhcp snooping limit rate <i>rate</i>	(Optional) Configure the number of DHCP packets per second that an interface can receive. The range is 1 to 2048. By default, no rate limit is configured. Note We recommend an untrusted rate limit of not more than 100 packets per second. If you configure rate limiting for trusted interfaces, you might need to increase the rate limit if the port is a trunk port assigned to more than one VLAN with DHCP snooping.
Step 9	exit	Return to global configuration mode.
Step 10	ip dhcp snooping verify mac-address	(Optional) Configure the switch to verify that the source MAC address in a DHCP packet received on untrusted ports matches the client hardware address in the packet. The default is to verify that the source MAC address matches the client hardware address in the packet.
Step 11	end	Return to privileged EXEC mode.
Step 12	show running-config	Verify your entries.
Step 13	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable DHCP snooping, use the **no ip dhcp snooping** global configuration command. To disable DHCP snooping on a VLAN or range of VLANs, use the **no ip dhcp snooping vlan** *vlan-range* global configuration command. To disable the insertion and the removal of the option-82 field, use the **no ip dhcp snooping information option** global configuration command. To configure an aggregation switch to drop incoming DHCP snooping packets with option-82 information from an edge switch, use the **no ip dhcp snooping information option allow-untrusted** global configuration command.

This example shows how to enable DHCP snooping globally and on VLAN 10 and to configure a rate limit of 100 packets per second on a port:

```
Switch(config)# ip dhcp snooping
Switch(config)# ip dhcp snooping vlan 10
Switch(config)# ip dhcp snooping information option
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip dhcp snooping limit rate 100
```

Enabling the DHCP Snooping Binding Database Agent

Beginning in privileged EXEC mode, follow these steps to enable and configure the DHCP snooping binding database agent on the switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip dhcp snooping database { flash [<i>number</i>]:/ <i>filename</i> ftp :// <i>user:password@host/filename</i> http ://[<i>username:password</i>]@]{ <i>hostname</i> <i>host-ip</i> }/[<i>directory</i>]/ <i>image-name.tar</i> rtp :// <i>user@host/filename</i> } tftp :// <i>host/filename</i>	Specify the URL for the database agent or the binding file by using one of these forms: <ul style="list-style-type: none"> • flash[<i>number</i>]:/<i>filename</i> (Optional) Use the <i>number</i> parameter to specify the stack member number of the stack master. The range for <i>number</i> is 1 to 4. • ftp://<i>user:password@host/filename</i> • http://[<i>username:password</i>]@]{<i>hostname</i> <i>host-ip</i>}/[<i>directory</i>]/<i>image-name.tar</i> • rtp://<i>user@host/filename</i> • tftp://<i>host/filename</i>
Step 3	ip dhcp snooping database timeout <i>seconds</i>	Specify (in seconds) how long to wait for the database transfer process to finish before stopping the process. The default is 300 seconds. The range is 0 to 86400. Use 0 to define an infinite duration, which means to continue trying the transfer indefinitely.
Step 4	ip dhcp snooping database write-delay <i>seconds</i>	Specify the duration for which the transfer should be delayed after the binding database changes. The range is from 15 to 86400 seconds. The default is 300 seconds (5 minutes).
Step 5	end	Return to privileged EXEC mode.
Step 6	ip dhcp snooping binding <i>mac-address</i> vlan <i>vlan-id</i> ip-address interface <i>interface-id</i> expiry <i>seconds</i>	(Optional) Add binding entries to the DHCP snooping binding database. The <i>vlan-id</i> range is from 1 to 4904. The <i>seconds</i> range is from 1 to 4294967295. Enter this command for each entry that you add. Note Use this command when you are testing or debugging the switch.
Step 7	show ip dhcp snooping database [detail]	Display the status and statistics of the DHCP snooping binding database agent.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To stop using the database agent and binding files, use the **no ip dhcp snooping database** global configuration command. To reset the timeout or delay values, use the **ip dhcp snooping database timeout** *seconds* or the **ip dhcp snooping database write-delay** *seconds* global configuration command.

To clear the statistics of the DHCP snooping binding database agent, use the **clear ip dhcp snooping database statistics** privileged EXEC command. To renew the database, use the **renew ip dhcp snooping database** privileged EXEC command.

To delete binding entries from the DHCP snooping binding database, use the **no ip dhcp snooping binding** *mac-address* **vlan** *vlan-id* **ip-address** **interface** *interface-id* privileged EXEC command. Enter this command for each entry that you want to delete.

Displaying DHCP Snooping Information

To display the DHCP snooping information, use the privileged EXEC commands in [Table 1-2](#):

Table 1-2 Commands for Displaying DHCP Information

Command	Purpose
<code>show ip dhcp snooping</code>	Displays the DHCP snooping configuration for a switch
<code>show ip dhcp snooping binding</code>	Displays only the dynamically configured bindings in the DHCP snooping binding database, also referred to as a binding table.
<code>show ip dhcp snooping database</code>	Displays the DHCP snooping binding database status and statistics.
<code>show ip dhcp snooping statistics</code>	Displays the DHCP snooping statistics in summary or detail form.



Note If DHCP snooping is enabled and an interface changes to the down state, the switch does not delete the statically configured bindings.

Understanding IP Source Guard



Note To use the IP source guard feature, the switch must be running the LAN Base image.

IPSG is a security feature that restricts IP traffic on nonrouted, Layer 2 interfaces by filtering traffic based on the DHCP snooping binding database and on manually configured IP source bindings. You can use IP source guard to prevent traffic attacks if a host tries to use the IP address of its neighbor.

You can enable IP source guard when DHCP snooping is enabled on an untrusted interface. After IPSG is enabled on an interface, the switch blocks all IP traffic received on the interface except for DHCP packets allowed by DHCP snooping. A port access control list (ACL) is applied to the interface. The port ACL allows only IP traffic with a source IP address in the IP source binding table and denies all other traffic.



Note The port ACL takes precedence over any router ACLs or VLAN maps that affect the same interface.

The IP source binding table bindings are learned by DHCP snooping or are manually configured (static IP source bindings). An entry in this table has an IP address with its associated MAC address and VLAN number. The switch uses the IP source binding table only when IP source guard is enabled.

IPSG is supported only on Layer 2 ports, including access and trunk ports. You can configure IPSG with source IP address filtering or with source IP and MAC address filtering.

- [Source IP Address Filtering, page 1-14](#)
- [Source IP and MAC Address Filtering, page 1-14](#)
- [IP Source Guard for Static Hosts, page 1-14](#)

Source IP Address Filtering

When IPSG is enabled with this option, IP traffic is filtered based on the source IP address. The switch forwards IP traffic when the source IP address matches an entry in the DHCP snooping binding database or a binding in the IP source binding table.

When a DHCP snooping binding or static IP source binding is added, changed, or deleted on an interface, the switch modifies the port ACL by using the IP source binding changes and re-applies the port ACL to the interface.

If you enable IPSG on an interface on which IP source bindings (dynamically learned by DHCP snooping or manually configured) are not configured, the switch creates and applies a port ACL that denies all IP traffic on the interface. If you disable IP source guard, the switch removes the port ACL from the interface.

Source IP and MAC Address Filtering

IP traffic is filtered based on the source IP and MAC addresses. The switch forwards traffic only when the source IP and MAC addresses match an entry in the IP source binding table.

When address filtering is enabled, the switch filters IP and non-IP traffic. If the source MAC address of an IP or non-IP packet matches a valid IP source binding, the switch forwards the packet. The switch drops all other types of packets except DHCP packets.

The switch uses port security to filter source MAC addresses. The interface can shut down when a port-security violation occurs.

IP Source Guard for Static Hosts

**Note**

Do not use IPSG (IP source guard) for static hosts on uplink ports or trunk ports.

IPSG for static hosts extends the IPSG capability to non-DHCP and static environments. The previous IPSG used the entries created by DHCP snooping to validate the hosts connected to a switch. Any traffic received from a host without a valid DHCP binding entry is dropped. This security feature restricts IP traffic on nonrouted Layer 2 interfaces. It filters traffic based on the DHCP snooping binding database and on manually configured IP source bindings. The previous version of IPSG required a DHCP environment for IPSG to work.

IPSG for static hosts allows IPSG to work without DHCP. IPSG for static hosts relies on IP device tracking-table entries to install port ACLs. The switch creates static entries based on ARP requests or other IP packets to maintain the list of valid hosts for a given port. You can also specify the number of hosts allowed to send traffic to a given port. This is equivalent to port security at Layer 3.

IPSG for static hosts also supports dynamic hosts. If a dynamic host receives a DHCP-assigned IP address that is available in the IP DHCP snooping table, the same entry is learned by the IP device tracking table. In a stacked environment, when the master failover occurs, the IP source guard entries for static hosts attached to member ports are retained. When you enter the **show ip device tracking all EXEC** command, the IP device tracking table displays the entries as ACTIVE.

**Note**

Some IP hosts with multiple network interfaces can inject some invalid packets into a network interface. The invalid packets contain the IP or MAC address for another network interface of the host as the source address. The invalid packets can cause IPSG for static hosts to connect to the host, to learn the invalid IP or MAC address bindings, and to reject the valid bindings. Consult the vendor of the corresponding operating system and the network interface to prevent the host from injecting invalid packets.

IPSG for static hosts initially learns IP or MAC bindings dynamically through an ACL-based snooping mechanism. IP or MAC bindings are learned from static hosts by ARP and IP packets. They are stored in the device tracking database. When the number of IP addresses that have been dynamically learned or statically configured on a given port reaches a maximum, the hardware drops any packet with a new IP address. To resolve hosts that have moved or gone away for any reason, IPSG for static hosts leverages IP device tracking to age out dynamically learned IP address bindings. This feature can be used with DHCP snooping. Multiple bindings are established on a port that is connected to both DHCP and static hosts. For example, bindings are stored in both the device tracking database as well as in the DHCP snooping binding database.

Configuring IP Source Guard

- [Default IP Source Guard Configuration, page 1-15](#)
- [IP Source Guard Configuration Guidelines, page 1-15](#)
- [Enabling IP Source Guard, page 1-16](#)
- [Configuring IP Source Guard for Static Hosts, page 1-17](#)

Default IP Source Guard Configuration

By default, IP source guard is disabled.

IP Source Guard Configuration Guidelines

- You can configure static IP bindings only on nonrouted ports. If you enter the **ip source binding mac-address vlan vlan-id ip-address interface interface-id** global configuration command on a routed interface, this error message appears:

```
Static IP source binding can only be configured on switch port.
```
- When IP source guard with source IP filtering is enabled on an interface, DHCP snooping must be enabled on the access VLAN for that interface.
- If you are enabling IP source guard on a trunk interface with multiple VLANs and DHCP snooping is enabled on all the VLANs, the source IP address filter is applied on all the VLANs.

**Note**

If IP source guard is enabled and you enable or disable DHCP snooping on a VLAN on the trunk interface, the switch might not properly filter traffic.

- If you enable IP source guard with source IP and MAC address filtering, DHCP snooping and port security must be enabled on the interface. You must also enter the **ip dhcp snooping information option** global configuration command and ensure that the DHCP server supports option 82. When IP source guard is enabled with MAC address filtering, the DHCP host MAC address is not learned until the host is granted a lease. When forwarding packets from the server to the host, DHCP snooping uses option-82 data to identify the host port.
- When configuring IP source guard on interfaces on which a private VLAN is configured, port security is not supported.
- IP source guard is not supported on EtherChannels.
- You can enable this feature when 802.1x port-based authentication is enabled.
- If the number of ternary content addressable memory (TCAM) entries exceeds the maximum, the CPU usage increases.
- In a switch stack, if IP source guard is configured on a stack member interface and you remove the switch configuration by entering the **no switch stack-member-number provision** global configuration command, the interface static bindings are removed from the binding table. They are not removed from the running configuration. If you again provision the switch by entering the **switch stack-member-number provision** command, the binding is restored. To remove the binding from the running configuration, you must disable IP source guard before entering the **no switch provision** global configuration command. The configuration is also removed if the switch reloads while the interface is removed from the binding table. For more information about provisioned switches, see the “[Stack Offline Configuration](#)” section on page 1-7.

Enabling IP Source Guard

Begin in privileged EXEC mode.

	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 verify source or ip verify source port-security	Enable IP source guard with source IP address filtering. Enable IP source guard with source IP and MAC address filtering. When you enable both IP source guard and Port Security by using the ip verify source port-security interface configuration command, there are two caveats: <ul style="list-style-type: none"> • The DHCP server must support option 82, or the client is not assigned an IP address. • The MAC address in the DHCP packet is not learned as a secure address. The MAC address of the DHCP client is learned as a secure address only when the switch receives non-DHCP data traffic.
Step 4	exit	Return to global configuration mode.
Step 5	ip source binding <i>mac-address</i> vlan <i>vlan-id</i> <i>ip-address</i> interface <i>interface-id</i>	Add a static IP source binding. Enter this command for each static binding.
Step 6	end	Return to privileged EXEC mode.

	Command	Purpose
Step 7	show ip verify source [<i>interface interface-id</i>]	Verify the IP source guard configuration.
Step 8	show ip source binding [<i>ip-address</i>] [<i>mac-address</i>] [dhcp-snooping static] [<i>interface interface-id</i>] [vlan <i>vlan-id</i>]	Display the IP source bindings on the switch, on a specific VLAN, or on a specific interface.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable IP source guard with source IP address filtering, use the **no ip verify source** interface configuration command.

To delete a static IP source binding entry, use the **no ip source** global configuration command.

This example shows how to enable IP source guard with source IP and MAC filtering on VLANs 10 and 11:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip verify source port-security
Switch(config-if)# exit
Switch(config)# ip source binding 0100.0022.0010 vlan 10 10.0.0.2 interface
gigabitethernet0/1
Switch(config)# ip source binding 0100.0230.0002 vlan 11 10.0.0.4 interface
gigabitethernet0/1
Switch(config)# end
```

Configuring IP Source Guard for Static Hosts

Configuring IP Source Guard for Static Hosts on a Layer 2 Access Port



Note You must configure the **ip device tracking maximum** *limit-number* interface configuration command globally for IPSG for static hosts to work. If you only configure this command on a port without enabling IP device tracking globally or by setting an IP device tracking maximum on that interface, IPSG with static hosts rejects all the IP traffic from that interface

Beginning in privileged EXEC mode:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip device tracking	Turn on the IP host table, and globally enable IP device tracking.
Step 3	interface <i>interface-id</i>	Enter interface configuration mode.
Step 4	switchport mode access	Configure a port as access.
Step 5	switchport access vlan <i>vlan-id</i>	Configure the VLAN for this port.

	Command	Purpose
Step 6	ip verify source tracking port-security	Enable IPSG for static hosts with MAC address filtering. Note When you enable both IP source guard and port security by using the ip verify source port-security interface configuration command: <ul style="list-style-type: none"> • The DHCP server must support option 82, or the client is not assigned an IP address. • The MAC address in the DHCP packet is not learned as a secure address. The MAC address of the DHCP client is learned as a secure address only when the switch receives non-DHCP data traffic.
Step 7	ip device tracking maximum <i>number</i>	Establish a maximum limit for the number of static IPs that the IP device tracking table allows on the port. The range is 1 to 10. The maximum number is 10. Note You must configure the ip device tracking maximum limit-number interface configuration command.
Step 8	switchport port-security	(Optional) Activate port security for this port.
Step 9	switchport port-security maximum <i>value</i>	(Optional) Establish a maximum of MAC addresses for this port.
Step 10	end	Return to privileged EXEC mode.
Step 11	show ip verify source interface <i>interface-id</i>	Verify the configuration and display IPSG permit ACLs for static hosts.
Step 12	show ip device track all [active inactive] count	Verify the configuration by displaying the IP-to-MAC binding for a given host on the switch interface. <ul style="list-style-type: none"> • all active—display only the active IP or MAC binding entries • all inactive—display only the inactive IP or MAC binding entries • all—display the active and inactive IP or MAC binding entries

This example shows how to stop IPSG with static hosts on an interface.

```
Switch(config-if)# no ip verify source
Switch(config-if)# no ip device tracking max
```

This example shows how to enable IPSG with static hosts on a port.

```
Switch(config)# ip device tracking
Switch(config)# ip device tracking max 10
Switch(config-if)# ip verify source tracking port-security
```

This example shows how to enable IPSG for static hosts with IP filters on a Layer 2 access port and to verify the valid IP bindings on the interface Gi0/3:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip device tracking
Switch(config)# interface gigabitethernet 0/3
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan 10
Switch(config-if)# ip device tracking maximum 5
Switch(config-if)# ip verify source tracking
Switch(config-if)# end

Switch# show ip verify source
Interface  Filter-type  Filter-mode  IP-address      Mac-address      Vlan
-----  -
Gi0/3     ip trk       active       40.1.1.24      40.1.1.24      10
Gi0/3     ip trk       active       40.1.1.20      40.1.1.20      10
Gi0/3     ip trk       active       40.1.1.21      40.1.1.21      10
```

This example shows how to enable IPSG for static hosts with IP-MAC filters on a Layer 2 access port, to verify the valid IP-MAC bindings on the interface Gi0/3, and to verify that the number of bindings on this interface has reached the maximum:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip device tracking
Switch(config)# interface gigabitethernet 0/3
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan 1
Switch(config-if)# ip device tracking maximum 5
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security maximum 5
Switch(config-if)# ip verify source tracking port-security
Switch(config-if)# end

Switch# show ip verify source
Interface  Filter-type  Filter-mode  IP-address      Mac-address      Vlan
-----  -
Gi0/3     ip-mac trk  active       40.1.1.24      00:00:00:00:03:04  1
Gi0/3     ip-mac trk  active       40.1.1.20      00:00:00:00:03:05  1
Gi0/3     ip-mac trk  active       40.1.1.21      00:00:00:00:03:06  1
Gi0/3     ip-mac trk  active       40.1.1.22      00:00:00:00:03:07  1
Gi0/3     ip-mac trk  active       40.1.1.23      00:00:00:00:03:08  1
```

This example displays all IP or MAC binding entries for all interfaces. The CLI displays all active as well as inactive entries. When a host is learned on a interface, the new entry is marked as active. When the same host is disconnected from that interface and connected to a different interface, a new IP or MAC binding entry displays as active as soon as the host is detected. The old entry for this host on the previous interface is marked as INACTIVE.

```
Switch# show ip device tracking all
IP Device Tracking = Enabled
IP Device Tracking Probe Count = 3
IP Device Tracking Probe Interval = 30
-----
   IP Address      MAC Address      Vlan  Interface      STATE
-----
200.1.1.8         0001.0600.0000  8     GigabitEthernet0/1  INACTIVE
200.1.1.9         0001.0600.0000  8     GigabitEthernet0/1  INACTIVE
200.1.1.10        0001.0600.0000  8     GigabitEthernet0/1  INACTIVE
200.1.1.1         0001.0600.0000  9     GigabitEthernet0/2  ACTIVE
200.1.1.1         0001.0600.0000  8     GigabitEthernet0/1  INACTIVE
```

```

200.1.1.2      0001.0600.0000  9  GigabitEthernet0/2  ACTIVE
200.1.1.2      0001.0600.0000  8  GigabitEthernet0/1  INACTIVE
200.1.1.3      0001.0600.0000  9  GigabitEthernet0/2  ACTIVE
200.1.1.3      0001.0600.0000  8  GigabitEthernet0/1  INACTIVE
200.1.1.4      0001.0600.0000  9  GigabitEthernet0/2  ACTIVE
200.1.1.4      0001.0600.0000  8  GigabitEthernet0/1  INACTIVE
200.1.1.5      0001.0600.0000  9  GigabitEthernet0/2  ACTIVE
200.1.1.5      0001.0600.0000  8  GigabitEthernet0/1  INACTIVE
200.1.1.6      0001.0600.0000  8  GigabitEthernet0/1  INACTIVE
200.1.1.7      0001.0600.0000  8  GigabitEthernet0/1  INACTIVE

```

This example displays all active IP or MAC binding entries for all interfaces:

```

Switch# show ip device tracking all active
IP Device Tracking = Enabled
IP Device Tracking Probe Count = 3
IP Device Tracking Probe Interval = 30
-----
   IP Address      MAC Address      Vlan  Interface          STATE
-----
200.1.1.1         0001.0600.0000  9     GigabitEthernet0/1  ACTIVE
200.1.1.2         0001.0600.0000  9     GigabitEthernet0/1  ACTIVE
200.1.1.3         0001.0600.0000  9     GigabitEthernet0/1  ACTIVE
200.1.1.4         0001.0600.0000  9     GigabitEthernet0/1  ACTIVE
200.1.1.5         0001.0600.0000  9     GigabitEthernet0/1  ACTIVE

```

This example displays all inactive IP or MAC binding entries for all interfaces. The host was first learned on GigabitEthernet 0/1 and then moved to GigabitEthernet 0/2. the IP or MAC binding entries learned on GigabitEthernet 0/1 are marked as inactive.

```

Switch# show ip device tracking all inactive
IP Device Tracking = Enabled
IP Device Tracking Probe Count = 3
IP Device Tracking Probe Interval = 30
-----
   IP Address      MAC Address      Vlan  Interface          STATE
-----
200.1.1.8         0001.0600.0000  8     GigabitEthernet0/1  INACTIVE
200.1.1.9         0001.0600.0000  8     GigabitEthernet0/1  INACTIVE
200.1.1.10        0001.0600.0000  8     GigabitEthernet0/1  INACTIVE
200.1.1.1         0001.0600.0000  8     GigabitEthernet0/1  INACTIVE
200.1.1.2         0001.0600.0000  8     GigabitEthernet0/1  INACTIVE
200.1.1.3         0001.0600.0000  8     GigabitEthernet0/1  INACTIVE
200.1.1.4         0001.0600.0000  8     GigabitEthernet0/1  INACTIVE
200.1.1.5         0001.0600.0000  8     GigabitEthernet0/1  INACTIVE
200.1.1.6         0001.0600.0000  8     GigabitEthernet0/1  INACTIVE
200.1.1.7         0001.0600.0000  8     GigabitEthernet0/1  INACTIVE

```

This example displays the count of all IP device tracking host entries for all interfaces:

```

Switch# show ip device tracking all count
Total IP Device Tracking Host entries: 5
-----
   Interface          Maximum Limit      Number of Entries
-----
Gi0/3                 5

```

Displaying IP Source Guard Information

To display the IP source guard information, use one or more of the privileged EXEC commands in [Table 1-3](#):

Table 1-3 *Commands for Displaying IP Source Guard Information*

Command	Purpose
<code>show ip device tracking</code>	Display the active IP or MAC binding entries for all interfaces.
<code>show ip source binding</code>	Display the IP source bindings on a switch.
<code>show ip verify source</code>	Display the IP source guard configuration on the switch.

Understanding DHCP Server Port-Based Address Allocation

DHCP server port-based address allocation is a feature that enables DHCP to maintain the same IP address on an Ethernet switch port regardless of the attached device client identifier or client hardware address.

When Ethernet switches are deployed in the network, they offer connectivity to the directly connected devices. In some environments, such as on a factory floor, if a device fails, the replacement device must be working immediately in the existing network. With the current DHCP implementation, there is no guarantee that DHCP would offer the same IP address to the replacement device. Control, monitoring, and other software expect a stable IP address associated with each device. If a device is replaced, the address assignment should remain stable even though the DHCP client has changed.

When configured, the DHCP server port-based address allocation feature ensures that the same IP address is always offered to the same connected port even as the client identifier or client hardware address changes in the DHCP messages received on that port. The DHCP protocol recognizes DHCP clients by the client identifier option in the DHCP packet. Clients that do not include the client identifier option are identified by the client hardware address. When you configure this feature, the port name of the interface overrides the client identifier or hardware address and the actual point of connection, the switch port, becomes the client identifier.

In all cases, by connecting the Ethernet cable to the same port, the same IP address is allocated through DHCP to the attached device.

The DHCP server port-based address allocation feature is only supported on a Cisco IOS DHCP server and not a third-party server.

Configuring DHCP Server Port-Based Address Allocation

- [Default Port-Based Address Allocation Configuration, page 1-21](#)
- [Port-Based Address Allocation Configuration Guidelines, page 1-22](#)
- [Enabling DHCP Server Port-Based Address Allocation, page 1-22](#)

Default Port-Based Address Allocation Configuration

By default, DHCP server port-based address allocation is disabled.

Port-Based Address Allocation Configuration Guidelines

These are the configuration guidelines for DHCP port-based address allocation:

- Only one IP address can be assigned per port.
- Reserved addresses (preassigned) cannot be cleared by using the **clear ip dhcp binding** global configuration command.
- Preassigned addresses are automatically excluded from normal dynamic IP address assignment. Preassigned addresses cannot be used in host pools, but there can be multiple preassigned addresses per DHCP address pool.
- To restrict assignments from the DHCP pool to preconfigured reservations (unreserved addresses are not offered to the client and other clients are not served by the pool), you can enter the **reserved-only** DHCP pool configuration command.

Enabling DHCP Server Port-Based Address Allocation

Beginning in privileged EXEC mode, follow these steps to globally enable port-based address allocation and to automatically generate a subscriber identifier on an interface.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip dhcp use subscriber-id client-id	Configure the DHCP server to globally use the subscriber identifier as the client identifier on all incoming DHCP messages.
Step 3	ip dhcp subscriber-id interface-name	Automatically generate a subscriber identifier based on the short name of the interface. A subscriber identifier configured on a specific interface takes precedence over this command.
Step 4	interface interface-id	Specify the interface to be configured, and enter interface configuration mode.
Step 5	ip dhcp server use subscriber-id client-id	Configure the DHCP server to use the subscriber identifier as the client identifier on all incoming DHCP messages on the interface.
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.

After enabling DHCP port-based address allocation on the switch, use the **ip dhcp pool** global configuration command to preassign IP addresses and to associate them to clients. To restrict assignments from the DHCP pool to preconfigured reservations, you can enter the **reserved-only** DHCP pool configuration command. Unreserved addresses that are part of the network or on pool ranges are not offered to the client, and other clients are not served by the pool. By entering this command, users can configure a group of switches with DHCP pools that share a common IP subnet and that ignore requests from clients of other switches.

Beginning in privileged EXEC mode follow these steps to preassign an IP address and to associate it to a client identified by the interface name.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip dhcp pool <i>poolname</i>	Enter DHCP pool configuration mode, and define the name for the DHCP pool. The pool name can be a symbolic string (such as Engineering) or an integer (such as 0).
Step 3	network <i>network-number</i> [<i>mask</i> <i>/prefix-length</i>]	Specify the subnet network number and mask of the DHCP address pool.
Step 4	address <i>ip-address</i> client-id <i>string</i> [ascii]	Reserve an IP address for a DHCP client identified by the interface name. <i>string</i> —can be an ASCII value or a hexadecimal value.
Step 5	reserved-only	(Optional) Use only reserved addresses in the DHCP address pool. The default is to not restrict pool addresses.
Step 6	end	Return to privileged EXEC mode.
Step 7	show ip dhcp pool	Verify DHCP pool configuration.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable DHCP port-based address allocation, use the **no ip dhcp use subscriber-id client-id** global configuration command. To disable the automatic generation of a subscriber identifier, use the **no ip dhcp subscriber-id interface-name** global configuration command. To disable the subscriber identifier on an interface, use the **no ip dhcp server use subscriber-id client-id** interface configuration command.

To remove an IP address reservation from a DHCP pool, use the **no address ip-address client-id string** DHCP pool configuration command. To change the address pool to nonrestricted, enter the **no reserved-only** DHCP pool configuration command.

In this example, a subscriber identifier is automatically generated, and the DHCP server ignores any client identifier fields in the DHCP messages and uses the subscriber identifier instead. The subscriber identifier is based on the short name of the interface and the client preassigned IP address 10.1.1.7.

```
Switch# show running config
Building configuration...
Current configuration : 4899 bytes
!
version 12.2
!
hostname switch
!
no aaa new-model
clock timezone EST 0
ip subnet-zero
ip dhcp relay information policy removal pad
no ip dhcp use vrf connected
ip dhcp use subscriber-id client-id
ip dhcp subscriber-id interface-name
ip dhcp excluded-address 10.1.1.1 10.1.1.3
!
```

```
ip dhcp pool dhcpool
network 10.1.1.0 255.255.255.0
address 10.1.1.7 client-id "Et1/0" ascii
<output truncated>
```

This example shows that the preassigned address was correctly reserved in the DHCP pool:

```
Switch# show ip dhcp pool dhcpool
Pool dhcp pool:
Utilization mark (high/low) : 100 / 0
Subnet size (first/next) : 0 / 0
Total addresses : 254
Leased addresses : 0
Excluded addresses : 4
Pending event : none
1 subnet is currently in the pool:
Current index   IP address range      Leased/Excluded/Total
10.1.1.1       10.1.1.1 - 10.1.1.254  0 / 4 / 254
1 reserved address is currently in the pool
Address        Client
10.1.1.7      Et1/0
```

For more information about configuring the DHCP server port-based address allocation feature, go to Cisco.com, and enter *Cisco IOS IP Addressing Services* in the Search field to access the Cisco IOS software documentation. You can also access the documentation:

http://www.cisco.com/en/US/docs/ios/ipaddr/command/reference/iad_book.html

Displaying DHCP Server Port-Based Address Allocation

To display the DHCP server port-based address allocation information, use one or more of the privileged EXEC commands in [Table 1-4](#):

Table 1-4 Commands for Displaying DHCP Port-Based Address Allocation Information

Command	Purpose
<code>show interface interface id</code>	Display the status and configuration of a specific interface.
<code>show ip dhcp pool</code>	Display the DHCP address pools.
<code>show ip dhcp binding</code>	Display address bindings on the Cisco IOS DHCP server.



CHAPTER 1

Configuring Dynamic ARP Inspection

This chapter describes how to configure dynamic Address Resolution Protocol inspection (dynamic ARP inspection) on the Catalyst 2960 and 2960-S switches. This feature helps prevent malicious attacks on the switch by not relaying invalid ARP requests and responses to other ports in the same VLAN. To use Dynamic ARP inspection, the switch must be running the LAN Base image.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.



Note


For complete syntax and usage information for the commands used in this chapter, see the command reference for this release.

- [Understanding Dynamic ARP Inspection, page 1-1](#)
- [Configuring Dynamic ARP Inspection, page 1-5](#)
- [Displaying Dynamic ARP Inspection Information, page 1-15](#)

Understanding Dynamic ARP Inspection

ARP provides IP communication within a Layer 2 broadcast domain by mapping an IP address to a MAC address. For example, Host B wants to send information to Host A but does not have the MAC address of Host A in its ARP cache. Host B generates a broadcast message for all hosts within the broadcast domain to obtain the MAC address associated with the IP address of Host A. All hosts within the broadcast domain receive the ARP request, and Host A responds with its MAC address. However, because ARP allows a gratuitous reply from a host even if an ARP request was not received, an ARP spoofing attack and the poisoning of ARP caches can occur. After the attack, all traffic from the device under attack flows through the attacker's computer and then to the router, switch, or host.

A malicious user can attack hosts, switches, and routers connected to your Layer 2 network by poisoning the ARP caches of systems connected to the subnet and by intercepting traffic intended for other hosts on the subnet. [Figure 1-1](#) shows an example of ARP cache poisoning.



Hosts A, B, and C are connected to the switch on interfaces A, B and C, all of which are on the same subnet. Their IP and MAC addresses are shown in parentheses; for example, Host A uses IP address IA and MAC address MA. When Host A needs to communicate to Host B at the IP layer, it broadcasts an ARP request for the MAC address associated with IP address IB. When the switch and Host B receive the ARP request, they populate their ARP caches with an ARP binding for a host with the IP address IA and a MAC address MA; for example, IP address IA is bound to MAC address MA. When Host B responds, the switch and Host A populate their ARP caches with a binding for a host with the IP address IB and the MAC address MB.

Host C can poison the ARP caches of the switch, Host A, and Host B by broadcasting forged ARP responses with bindings for a host with an IP address of IA (or IB) and a MAC address of MC. Hosts with poisoned ARP caches use the MAC address MC as the destination MAC address for traffic intended for IA or IB. This means that Host C intercepts that traffic. Because Host C knows the true MAC addresses associated with IA and IB, it can forward the intercepted traffic to those hosts by using the correct MAC address as the destination. Host C has inserted itself into the traffic stream from Host A to Host B, the classic *man-in-the-middle* attack.

Dynamic ARP inspection is a security feature that validates ARP packets in a network. It intercepts, logs, and discards ARP packets with invalid IP-to-MAC address bindings. This capability protects the network from certain man-in-the-middle attacks.

Dynamic ARP inspection ensures that only valid ARP requests and responses are relayed. The switch performs these activities:

- Intercepts all ARP requests and responses on untrusted ports
- Verifies that each of these intercepted packets has a valid IP-to-MAC address binding before updating the local ARP cache or before forwarding the packet to the appropriate destination
- Drops invalid ARP packets

Dynamic ARP inspection determines the validity of an ARP packet based on valid IP-to-MAC address bindings stored in a trusted database, the DHCP snooping binding database. This database is built by DHCP snooping if DHCP snooping is enabled on the VLANs and on the switch. If the ARP packet is received on a trusted interface, the switch forwards the packet without any checks. On untrusted interfaces, the switch forwards the packet only if it is valid.

You enable dynamic ARP inspection on a per-VLAN basis by using the **ip arp inspection vlan *vlan-range*** global configuration command. For configuration information, see the “[Configuring Dynamic ARP Inspection in DHCP Environments](#)” section on page 1-7.

In non-DHCP environments, dynamic ARP inspection can validate ARP packets against user-configured ARP access control lists (ACLs) for hosts with statically configured IP addresses. You define an ARP ACL by using the **arp access-list *acl-name*** global configuration command. For configuration information, see the “[Configuring ARP ACLs for Non-DHCP Environments](#)” section on page 1-9. The switch logs dropped packets. For more information about the log buffer, see the “[Logging of Dropped Packets](#)” section on page 1-5.

You can configure dynamic ARP inspection to drop ARP packets when the IP addresses in the packets are invalid or when the MAC addresses in the body of the ARP packets do not match the addresses specified in the Ethernet header. Use the **ip arp inspection validate** {[src-mac] [dst-mac] [ip]} global configuration command. For more information, see the “Performing Validation Checks” section on page 1-12.

Interface Trust States and Network Security

Dynamic ARP inspection associates a trust state with each interface on the switch. Packets arriving on trusted interfaces bypass all dynamic ARP inspection validation checks, and those arriving on untrusted interfaces undergo the dynamic ARP inspection validation process.

In a typical network configuration, you configure all switch ports connected to host ports as untrusted and configure all switch ports connected to switches as trusted. With this configuration, all ARP packets entering the network from a given switch bypass the security check. No other validation is needed at any other place in the VLAN or in the network. You configure the trust setting by using the **ip arp inspection trust** interface configuration command.




Caution

Use the trust state configuration carefully. Configuring interfaces as untrusted when they should be trusted can result in a loss of connectivity.

In [Figure 1-2](#), assume that both Switch A and Switch B are running dynamic ARP inspection on the VLAN that includes Host 1 and Host 2. If Host 1 and Host 2 acquire their IP addresses from the DHCP server connected to Switch A, only Switch A binds the IP-to-MAC address of Host 1. Therefore, if the interface between Switch A and Switch B is untrusted, the ARP packets from Host 1 are dropped by Switch B. Connectivity between Host 1 and Host 2 is lost.

Figure 1-2 ARP Packet Validation on a VLAN Enabled for Dynamic ARP Inspection



Configuring interfaces to be trusted when they are actually untrusted leaves a security hole in the network. If Switch A is not running dynamic ARP inspection, Host 1 can easily poison the ARP cache of Switch B (and Host 2, if the link between the switches is configured as trusted). This condition can occur even though Switch B is running dynamic ARP inspection.

Dynamic ARP inspection ensures that hosts (on untrusted interfaces) connected to a switch running dynamic ARP inspection do not poison the ARP caches of other hosts in the network. However, dynamic ARP inspection does not prevent hosts in other portions of the network from poisoning the caches of the hosts that are connected to a switch running dynamic ARP inspection.

In cases in which some switches in a VLAN run dynamic ARP inspection and other switches do not, configure the interfaces connecting such switches as untrusted. However, to validate the bindings of packets from nondynamic ARP inspection switches, configure the switch running dynamic ARP inspection with ARP ACLs. When you cannot determine such bindings, at Layer 3, isolate switches running dynamic ARP inspection from switches not running dynamic ARP inspection switches. For configuration information, see the [“Configuring ARP ACLs for Non-DHCP Environments”](#) section on page 1-9.



Note

Depending on the setup of the DHCP server and the network, it might not be possible to validate a given ARP packet on all switches in the VLAN.

Rate Limiting of ARP Packets

The switch CPU performs dynamic ARP inspection validation checks; therefore, the number of incoming ARP packets is rate-limited to prevent a denial-of-service attack. By default, the rate for untrusted interfaces is 15 packets per second (pps). Trusted interfaces are not rate-limited. You can change this setting by using the **ip arp inspection limit** interface configuration command.

When the rate of incoming ARP packets exceeds the configured limit, the switch places the port in the error-disabled state. The port remains in that state until you intervene. You can use the **errdisable recovery** global configuration command to enable error disable recovery so that ports automatically emerge from this state after a specified timeout period.



Note

The rate limit for an EtherChannel is applied separately to each switch in a stack. For example, if a limit of 20 pps is configured on the EtherChannel, each switch with ports in the EtherChannel can carry up to 20 pps. If any switch exceeds the limit, the entire EtherChannel is placed into the error-disabled state.

For configuration information, see the [“Limiting the Rate of Incoming ARP Packets”](#) section on page 1-11.

Relative Priority of ARP ACLs and DHCP Snooping Entries

Dynamic ARP inspection uses the DHCP snooping binding database for the list of valid IP-to-MAC address bindings.

ARP ACLs take precedence over entries in the DHCP snooping binding database. The switch uses ACLs only if you configure them by using the **ip arp inspection filter vlan** global configuration command. The switch first compares ARP packets to user-configured ARP ACLs. If the ARP ACL denies the ARP packet, the switch also denies the packet even if a valid binding exists in the database populated by DHCP snooping.

Logging of Dropped Packets

When the switch drops a packet, it places an entry in the log buffer and then generates system messages on a rate-controlled basis. After the message is generated, the switch clears the entry from the log buffer. Each log entry contains flow information, such as the receiving VLAN, the port number, the source and destination IP addresses, and the source and destination MAC addresses.

You use the **ip arp inspection log-buffer** global configuration command to configure the number of entries in the buffer and the number of entries needed in the specified interval to generate system messages. You specify the type of packets that are logged by using the **ip arp inspection vlan logging** global configuration command. For configuration information, see the [“Configuring the Log Buffer” section on page 1-13](#).

Configuring Dynamic ARP Inspection

- [Default Dynamic ARP Inspection Configuration, page 1-5](#)
- [Dynamic ARP Inspection Configuration Guidelines, page 1-6](#)
- [Configuring Dynamic ARP Inspection in DHCP Environments, page 1-7](#) (required in DHCP environments)
- [Configuring ARP ACLs for Non-DHCP Environments, page 1-9](#) (required in non-DHCP environments)
- [Limiting the Rate of Incoming ARP Packets, page 1-11](#) (optional)
- [Performing Validation Checks, page 1-12](#) (optional)
- [Configuring the Log Buffer, page 1-13](#) (optional)

Default Dynamic ARP Inspection Configuration

Table 1-1 shows the default dynamic ARP inspection configuration.

Table 1-1 Default Dynamic ARP Inspection Configuration

Feature	Default Setting
Dynamic ARP inspection	Disabled on all VLANs.
Interface trust state	All interfaces are untrusted.
Rate limit of incoming ARP packets	The rate is 15 pps on untrusted interfaces, assuming that the network is a switched network with a host connecting to as many as 15 new hosts per second. The rate is unlimited on all trusted interfaces. The burst interval is 1 second.
ARP ACLs for non-DHCP environments	No ARP ACLs are defined.
Validation checks	No checks are performed.

Table 1-1 Default Dynamic ARP Inspection Configuration (continued)

Feature	Default Setting
Log buffer	<p>When dynamic ARP inspection is enabled, all denied or dropped ARP packets are logged.</p> <p>The number of entries in the log is 32.</p> <p>The number of system messages is limited to 5 per second.</p> <p>The logging-rate interval is 1 second.</p>
Per-VLAN logging	All denied or dropped ARP packets are logged.

Dynamic ARP Inspection Configuration Guidelines

These are the dynamic ARP inspection configuration guidelines:

- Dynamic ARP inspection is an ingress security feature; it does not perform any egress checking.
- Dynamic ARP inspection is not effective for hosts connected to switches that do not support dynamic ARP inspection or that do not have this feature enabled. Because man-in-the-middle attacks are limited to a single Layer 2 broadcast domain, separate the domain with dynamic ARP inspection checks from the one with no checking. This action secures the ARP caches of hosts in the domain enabled for dynamic ARP inspection.
- Dynamic ARP inspection depends on the entries in the DHCP snooping binding database to verify IP-to-MAC address bindings in incoming ARP requests and ARP responses. Make sure to enable DHCP snooping to permit ARP packets that have dynamically assigned IP addresses. For configuration information, see [Chapter 1, “Configuring DHCP and IP Source Guard Features.”](#)
When DHCP snooping is disabled or in non-DHCP environments, use ARP ACLs to permit or to deny packets.
- Dynamic ARP inspection is supported on access ports, trunk ports, EtherChannel ports, and private VLAN ports.



Note Do not enable Dynamic ARP inspection on RSPAN VLANs. If Dynamic ARP inspection is enabled on RSPAN VLANs, Dynamic ARP inspection packets might not reach the RSPAN destination port.

- A physical port can join an EtherChannel port channel only when the trust state of the physical port and the channel port match. Otherwise, the physical port remains suspended in the port channel. A port channel inherits its trust state from the first physical port that joins the channel. Consequently, the trust state of the first physical port need not match the trust state of the channel.
Conversely, when you change the trust state on the port channel, the switch configures a new trust state on all the physical ports that comprise the channel.
- The rate limit is calculated separately on each switch in a switch stack. For a cross-stack EtherChannel, this means that the actual rate limit might be higher than the configured value. For example, if you set the rate limit to 30 pps on an EtherChannel that has one port on switch 1 and one port on switch 2, each port can receive packets at 29 pps without causing the EtherChannel to become error-disabled.

- The operating rate for the port channel is cumulative across all the physical ports within the channel. For example, if you configure the port channel with an ARP rate-limit of 400 pps, all the interfaces combined on the channel receive an aggregate 400 pps. The rate of incoming ARP packets on EtherChannel ports is equal to the sum of the incoming rate of packets from all the channel members. Configure the rate limit for EtherChannel ports only after examining the rate of incoming ARP packets on the channel-port members.

The rate of incoming packets on a physical port is checked against the port-channel configuration rather than the physical-ports configuration. The rate-limit configuration on a port channel is independent of the configuration on its physical ports.

If the EtherChannel receives more ARP packets than the configured rate, the channel (including all physical ports) is placed in the error-disabled state.

- Make sure to limit the rate of ARP packets on incoming trunk ports. Configure trunk ports with higher rates to reflect their aggregation and to handle packets across multiple dynamic ARP inspection-enabled VLANs. You also can use the **ip arp inspection limit none** interface configuration command to make the rate unlimited. A high rate-limit on one VLAN can cause a denial-of-service attack to other VLANs when the software places the port in the error-disabled state.
- When you enable dynamic ARP inspection on the switch, policers that were configured to police ARP traffic are no longer effective. The result is that all ARP traffic is sent to the CPU.

Configuring Dynamic ARP Inspection in DHCP Environments

This procedure shows how to configure dynamic ARP inspection when two switches support this feature. Host 1 is connected to Switch A, and Host 2 is connected to Switch B as shown in [Figure 1-2 on page 1-3](#). Both switches are running dynamic ARP inspection on VLAN 1 where the hosts are located. A DHCP server is connected to Switch A. Both hosts acquire their IP addresses from the same DHCP server. Therefore, Switch A has the bindings for Host 1 and Host 2, and Switch B has the binding for Host 2.



Note

Dynamic ARP inspection depends on the entries in the DHCP snooping binding database to verify IP-to-MAC address bindings in incoming ARP requests and ARP responses. Make sure to enable DHCP snooping to permit ARP packets that have dynamically assigned IP addresses. For configuration information, see [Chapter 1, “Configuring DHCP and IP Source Guard Features.”](#)

For information on how to configure dynamic ARP inspection when only one switch supports the feature, see the [“Configuring ARP ACLs for Non-DHCP Environments” section on page 1-9](#).

Beginning in privileged EXEC mode, follow these steps to configure dynamic ARP inspection. You must perform this procedure on both switches. This procedure is required.

	Command	Purpose
Step 1	show cdp neighbors	Verify the connection between the switches.
Step 2	configure terminal	Enter global configuration mode.
Step 3	ip arp inspection vlan <i>vlan-range</i>	Enable dynamic ARP inspection on a per-VLAN basis. By default, dynamic ARP inspection is disabled on all VLANs. For <i>vlan-range</i> , 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. Specify the same VLAN ID for both switches.
Step 4	interface <i>interface-id</i>	Specify the interface connected to the other switch, and enter interface configuration mode.
Step 5	ip arp inspection trust	Configure the connection between the switches as trusted. By default, all interfaces are untrusted. The switch does not check ARP packets that it receives from the other switch on the trusted interface. It simply forwards the packets. For untrusted interfaces, the switch intercepts all ARP requests and responses. It verifies that the intercepted packets have valid IP-to-MAC address bindings before updating the local cache and before forwarding the packet to the appropriate destination. The switch drops invalid packets and logs them in the log buffer according to the logging configuration specified with the ip arp inspection vlan logging global configuration command. For more information, see the “Configuring the Log Buffer” section on page 1-13.
Step 6	end	Return to privileged EXEC mode.
Step 7	show ip arp inspection interfaces show ip arp inspection vlan <i>vlan-range</i>	Verify the dynamic ARP inspection configuration.
Step 8	show ip dhcp snooping binding	Verify the DHCP bindings.
Step 9	show ip arp inspection statistics vlan <i>vlan-range</i>	Check the dynamic ARP inspection statistics.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable dynamic ARP inspection, use the **no ip arp inspection vlan** *vlan-range* global configuration command. To return the interfaces to an untrusted state, use the **no ip arp inspection trust** interface configuration command.

This example shows how to configure dynamic ARP inspection on Switch A in VLAN 1. You would perform a similar procedure on Switch B:

```
Switch(config)# ip arp inspection vlan 1
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip arp inspection trust
```

Configuring ARP ACLs for Non-DHCP Environments

This procedure shows how to configure dynamic ARP inspection when Switch B shown in [Figure 1-2 on page 1-3](#) does not support dynamic ARP inspection or DHCP snooping.

If you configure port 1 on Switch A as trusted, a security hole is created because both Switch A and Host 1 could be attacked by either Switch B or Host 2. To prevent this possibility, you must configure port 1 on Switch A as untrusted. To permit ARP packets from Host 2, you must set up an ARP ACL and apply it to VLAN 1. If the IP address of Host 2 is not static (it is impossible to apply the ACL configuration on Switch A) you must separate Switch A from Switch B at Layer 3 and use a router to route packets between them.

Beginning in privileged EXEC mode, follow these steps to configure an ARP ACL on Switch A. This procedure is required in non-DHCP environments.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	arp access-list <i>acl-name</i>	Define an ARP ACL, and enter ARP access-list configuration mode. By default, no ARP access lists are defined. Note At the end of the ARP access list, there is an implicit deny ip any mac any command.
Step 3	permit ip host <i>sender-ip</i> mac host <i>sender-mac</i> [log]	Permit ARP packets from the specified host (Host 2). <ul style="list-style-type: none"> For <i>sender-ip</i>, enter the IP address of Host 2. For <i>sender-mac</i>, enter the MAC address of Host 2. (Optional) Specify log to log a packet in the log buffer when it matches the access control entry (ACE). Matches are logged if you also configure the matchlog keyword in the ip arp inspection vlan logging global configuration command. For more information, see the “Configuring the Log Buffer” section on page 1-13.
Step 4	exit	Return to global configuration mode.

	Command	Purpose
Step 5	ip arp inspection filter <i>arp-acl-name</i> vlan <i>vlan-range</i> [static]	<p>Apply the ARP ACL to the VLAN. By default, no defined ARP ACLs are applied to any VLAN.</p> <ul style="list-style-type: none"> For <i>arp-acl-name</i>, specify the name of the ACL created in Step 2. For <i>vlan-range</i>, specify the VLAN that the switches and hosts are in. 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) Specify static to treat implicit denies in the ARP ACL as explicit denies and to drop packets that do not match any previous clauses in the ACL. DHCP bindings are not used. <p>If you do not specify this keyword, it means that there is no explicit deny in the ACL that denies the packet, and DHCP bindings determine whether a packet is permitted or denied if the packet does not match any clauses in the ACL.</p> <p>ARP packets containing only IP-to-MAC address bindings are compared against the ACL. Packets are permitted only if the access list permits them.</p>
Step 6	interface <i>interface-id</i>	Specify the Switch A interface that is connected to Switch B, and enter interface configuration mode.
Step 7	no ip arp inspection trust	<p>Configure the Switch A interface that is connected to Switch B as untrusted.</p> <p>By default, all interfaces are untrusted.</p> <p>For untrusted interfaces, the switch intercepts all ARP requests and responses. It verifies that the intercepted packets have valid IP-to-MAC address bindings before updating the local cache and before forwarding the packet to the appropriate destination. The switch drops invalid packets and logs them in the log buffer according to the logging configuration specified with the ip arp inspection vlan logging global configuration command. For more information, see the “Configuring the Log Buffer” section on page 1-13.</p>
Step 8	end	Return to privileged EXEC mode.
Step 9	show arp access-list [<i>acl-name</i>] show ip arp inspection vlan <i>vlan-range</i> show ip arp inspection interfaces	Verify your entries.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the ARP ACL, use the **no arp access-list** global configuration command. To remove the ARP ACL attached to a VLAN, use the **no ip arp inspection filter** *arp-acl-name* **vlan** *vlan-range* global configuration command.

This example shows how to configure an ARP ACL called *host2* on Switch A, to permit ARP packets from Host 2 (IP address 1.1.1.1 and MAC address 0001.0001.0001), to apply the ACL to VLAN 1, and to configure port 1 on Switch A as untrusted:

```
Switch(config)# arp access-list host2
Switch(config-arp-acl)# permit ip host 1.1.1.1 mac host 1.1.1
Switch(config-arp-acl)# exit
Switch(config)# ip arp inspection filter host2 vlan 1
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# no ip arp inspection trust
```

Limiting the Rate of Incoming ARP Packets

The switch CPU performs dynamic ARP inspection validation checks; therefore, the number of incoming ARP packets is rate-limited to prevent a denial-of-service attack.

When the rate of incoming ARP packets exceeds the configured limit, the switch places the port in the error-disabled state. The port remains in that state until you enable error-disabled recovery so that ports automatically emerge from this state after a specified timeout period.



Note

Unless you configure a rate limit on an interface, changing the trust state of the interface also changes its rate limit to the default value for that trust state. After you configure the rate limit, the interface retains the rate limit even when its trust state is changed. If you enter the **no ip arp inspection limit** interface configuration command, the interface reverts to its default rate limit.

For configuration guidelines for rate limiting trunk ports and EtherChannel ports, see the “[Dynamic ARP Inspection Configuration Guidelines](#)” section on page 1-6.

Beginning in privileged EXEC mode, follow these steps to limit the rate of incoming ARP packets. 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 rate-limited, and enter interface configuration mode.
Step 3	ip arp inspection limit { rate <i>pps</i> [burst interval <i>seconds</i>] none }	Limit the rate of incoming ARP requests and responses on the interface. The default rate is 15 pps on untrusted interfaces and unlimited on trusted interfaces. The burst interval is 1 second. The keywords have these meanings: <ul style="list-style-type: none"> For rate <i>pps</i>, specify an upper limit for the number of incoming packets processed per second. The range is 0 to 2048 pps. (Optional) For burst interval <i>seconds</i>, specify the consecutive interval in seconds, over which the interface is monitored for a high rate of ARP packets. The range is 1 to 15. For rate none, specify no upper limit for the rate of incoming ARP packets that can be processed.
Step 4	exit	Return to global configuration mode.

	Command	Purpose
Step 5	errdisable recovery cause arp-inspection interval <i>interval</i>	(Optional) Enable error recovery from the dynamic ARP inspection error-disable state. By default, recovery is disabled, and the recovery interval is 300 seconds. For interval <i>interval</i> , specify the time in seconds to recover from the error-disable state. The range is 30 to 86400.
Step 6	exit	Return to privileged EXEC mode.
Step 7	show ip arp inspection interfaces show errdisable recovery	Verify your settings.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default rate-limit configuration, use the **no ip arp inspection limit** interface configuration command. To disable error recovery for dynamic ARP inspection, use the **no errdisable recovery cause arp-inspection** global configuration command.

Performing Validation Checks

Dynamic ARP inspection intercepts, logs, and discards ARP packets with invalid IP-to-MAC address bindings. You can configure the switch to perform additional checks on the destination MAC address, the sender and target IP addresses, and the source MAC address.

Beginning in privileged EXEC mode, follow these steps to perform specific checks on incoming ARP packets. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip arp inspection validate {[src-mac] [dst-mac] [ip]}	<p>Perform a specific check on incoming ARP packets. By default, no checks are performed.</p> <p>The keywords have these meanings:</p> <ul style="list-style-type: none"> For src-mac, check the source MAC address in the Ethernet header against the sender MAC address in the ARP body. This check is performed on both ARP requests and responses. When enabled, packets with different MAC addresses are classified as invalid and are dropped. For dst-mac, check the destination MAC address in the Ethernet header against the target MAC address in ARP body. This check is performed for ARP responses. When enabled, packets with different MAC addresses are classified as invalid and are dropped. For ip, check the ARP body for invalid and unexpected IP addresses. Addresses include 0.0.0.0, 255.255.255.255, and all IP multicast addresses. Sender IP addresses are checked in all ARP requests and responses, and target IP addresses are checked only in ARP responses. <p>You must specify at least one of the keywords. Each command overrides the configuration of the previous command; that is, if a command enables src and dst mac validations, and a second command enables IP validation only, the src and dst mac validations are disabled as a result of the second command.</p>
Step 3	exit	Return to privileged EXEC mode.
Step 4	show ip arp inspection vlan <i>vlan-range</i>	Verify your settings.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable checking, use the **no ip arp inspection validate [src-mac] [dst-mac] [ip]** global configuration command. To display statistics for forwarded, dropped, and MAC and IP validation failure packets, use the **show ip arp inspection statistics** privileged EXEC command.

Configuring the Log Buffer

When the switch drops a packet, it places an entry in the log buffer and then generates system messages on a rate-controlled basis. After the message is generated, the switch clears the entry from the log buffer. Each log entry contains flow information, such as the receiving VLAN, the port number, the source and destination IP addresses, and the source and destination MAC addresses.

A log-buffer entry can represent more than one packet. For example, if an interface receives many packets on the same VLAN with the same ARP parameters, the switch combines the packets as one entry in the log buffer and generates a single system message for the entry.

If the log buffer overflows, it means that a log event does not fit into the log buffer, and the display for the **show ip arp inspection log** privileged EXEC command is affected. A -- in the display appears in place of all data except the packet count and the time. No other statistics are provided for the entry. If you see this entry in the display, increase the number of entries in the log buffer or increase the logging rate.

The log buffer configuration applies to each stack member in a switch stack. Each stack member has the specified **logs number** entries and generates system messages at the configured rate. For example, if the interval (rate) is one entry per second, up to five system messages are generated per second in a five-member switch stack.

Beginning in privileged EXEC mode, follow these steps to configure the log buffer. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip arp inspection log-buffer {entries number logs number interval seconds}	<p>Configure the dynamic ARP inspection logging buffer.</p> <p>By default, when dynamic ARP inspection is enabled, denied or dropped ARP packets are logged. The number of log entries is 32. The number of system messages is limited to 5 per second. The logging-rate interval is 1 second.</p> <p>The keywords have these meanings:</p> <ul style="list-style-type: none"> • For entries number, specify the number of entries to be logged in the buffer. The range is 0 to 1024. • For logs number interval seconds, specify the number of entries to generate system messages in the specified interval. <p>For logs number, the range is 0 to 1024. A 0 value means that the entry is placed in the log buffer, but a system message is not generated.</p> <p>For interval seconds, the range is 0 to 86400 seconds (1 day). A 0 value means that a system message is immediately generated (and the log buffer is always empty).</p> <p>An interval setting of 0 overrides a log setting of 0.</p> <p>The logs and interval settings interact. If the logs number X is greater than interval seconds Y, X divided by Y (X/Y) system messages are sent every second. Otherwise, one system message is sent every Y divided by X (Y/X) seconds.</p>

	Command	Purpose
Step 3	ip arp inspection vlan <i>vlan-range</i> logging { acl-match { matchlog none } dhcp-bindings { all none permit }}	Control the type of packets that are logged per VLAN. By default, all denied or all dropped packets are logged. The term <i>logged</i> means the entry is placed in the log buffer and a system message is generated. The keywords have these meanings: <ul style="list-style-type: none"> For <i>vlan-range</i>, 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 acl-match matchlog, log packets based on the ACE logging configuration. If you specify the matchlog keyword in this command and the log keyword in the permit or deny ARP access-list configuration command, ARP packets permitted or denied by the ACL are logged. For acl-match none, do not log packets that match ACLs. For dhcp-bindings all, log all packets that match DHCP bindings. For dhcp-bindings none, do not log packets that match DHCP bindings. For dhcp-bindings permit, log DHCP-binding permitted packets.
Step 4	exit	Return to privileged EXEC mode.
Step 5	show ip arp inspection log	Verify your settings.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default log buffer settings, use the **no ip arp inspection log-buffer** {**entries** | **logs**} global configuration command. To return to the default VLAN log settings, use the **no ip arp inspection vlan** *vlan-range* **logging** {**acl-match** | **dhcp-bindings**} global configuration command. To clear the log buffer, use the **clear ip arp inspection log** privileged EXEC command.

Displaying Dynamic ARP Inspection Information

To display dynamic ARP inspection information, use the privileged EXEC commands described in [Table 1-2](#):

Table 1-2 Commands for Displaying Dynamic ARP Inspection Information

Command	Description
show arp access-list [<i>acl-name</i>]	Displays detailed information about ARP ACLs.
show ip arp inspection interfaces [<i>interface-id</i>]	Displays the trust state and the rate limit of ARP packets for the specified interface or all interfaces.
show ip arp inspection vlan <i>vlan-range</i>	Displays the configuration and the operating state of dynamic ARP inspection for the specified VLAN. If no VLANs are specified or if a range is specified, displays information only for VLANs with dynamic ARP inspection enabled (active).

To clear or display dynamic ARP inspection statistics, use the privileged EXEC commands in [Table 1-3](#):

Table 1-3 *Commands for Clearing or Displaying Dynamic ARP Inspection Statistics*

Command	Description
clear ip arp inspection statistics	Clears dynamic ARP inspection statistics.
show ip arp inspection statistics [vlan <i>vlan-range</i>]	Displays statistics for forwarded, dropped, MAC validation failure, IP validation failure, ACL permitted and denied, and DHCP permitted and denied packets for the specified VLAN. If no VLANs are specified or if a range is specified, displays information only for VLANs with dynamic ARP inspection enabled (active).

For the **show ip arp inspection statistics** command, the switch increments the number of forwarded packets for each ARP request and response packet on a trusted dynamic ARP inspection port. The switch increments the number of ACL or DHCP permitted packets for each packet that is denied by source MAC, destination MAC, or IP validation checks, and the switch increments the appropriate failure count.

To clear or display dynamic ARP inspection logging information, use the privileged EXEC commands in [Table 1-4](#):

Table 1-4 *Commands for Clearing or Displaying Dynamic ARP Inspection Logging Information*

Command	Description
clear ip arp inspection log	Clears the dynamic ARP inspection log buffer.
show ip arp inspection log	Displays the configuration and contents of the dynamic ARP inspection log buffer.

For more information about these commands, see the command reference for this release.



CHAPTER 1

Configuring IGMP Snooping and MVR



Note

To use MVR, the switch must be running the LAN Base image.

This chapter describes how to configure Internet Group Management Protocol (IGMP) snooping on the Catalyst 2960 and 2960-S switches, including an application of local IGMP snooping, Multicast VLAN Registration (MVR). It also includes procedures for controlling multicast group membership by using IGMP filtering and procedures for configuring the IGMP throttling action. Unless otherwise noted, the term *switch* refers to a standalone switch and a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.



Note

For complete syntax and usage information for the commands used in this chapter, see the switch command reference for this release and the “IP Multicast Routing Commands” section in the *Cisco IOS IP Command Reference, Volume 3 of 3: Multicast, Release 12.4* on Cisco.com.

This chapter consists of these sections:

- [Understanding IGMP Snooping, page 1-2](#)
- [Configuring IGMP Snooping, page 1-7](#)
- [Displaying IGMP Snooping Information, page 1-15](#)
- [Understanding Multicast VLAN Registration, page 1-17](#)
- [Configuring MVR, page 1-19](#)
- [Displaying MVR Information, page 1-22](#)
- [Configuring IGMP Filtering and Throttling, page 1-23](#)
- [Displaying IGMP Filtering and Throttling Configuration, page 1-28](#)



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, see RFC 1112 and RFC 2236.

The multicast router 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.

You can configure an IGMP snooping querier to support IGMP snooping in subnets without multicast interfaces because the multicast traffic does not need to be routed. For more information about the IGMP snooping querier, see the “[Configuring the IGMP Snooping Querier](#)” section on page 1-14.

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 IGMP snooping characteristics:

- [IGMP Versions, page 1-3](#)
- [Joining a Multicast Group, page 1-3](#)
- [Leaving a Multicast Group, page 1-5](#)
- [Immediate Leave, page 1-5](#)
- [IGMP Configurable-Leave Timer, page 1-6](#)
- [IGMP Report Suppression, page 1-6](#)

IGMP Versions

The switch supports IGMP Version 1, IGMP Version 2, and IGMP Version 3. These versions are interoperable on the switch. For example, if IGMP snooping is enabled on an IGMPv2 switch and the switch receives an IGMPv3 report from a host, the switch can forward the IGMPv3 report to the multicast router.

**Note**

The switch supports IGMPv3 snooping based only on the destination multicast MAC address. It does not support snooping based on the source MAC address or on proxy reports.

An IGMPv3 switch supports Basic IGMPv3 Snooping Support (BISS), which includes support for the snooping features on IGMPv1 and IGMPv2 switches and for IGMPv3 membership report messages. BISS constrains the flooding of multicast traffic when your network includes IGMPv3 hosts. It constrains traffic to approximately the same set of ports as the IGMP snooping feature on IGMPv2 or IGMPv1 hosts.

**Note**

IGMPv3 join and leave messages are not supported on switches running IGMP filtering or MVR.

An IGMPv3 switch can receive messages from and forward messages to a device running the Source Specific Multicast (SSM) feature. For more information about source-specific multicast with IGMPv3 and IGMP:


http://www.cisco.com/en/US/docs/ios/12_1t/12_1t5/feature/guide/dtssm5t.html

Joining a Multicast Group

When a host connected to the switch wants to join an IP multicast group and 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 1-1](#).

Figure 1-1 Initial IGMP Join Message



Router A sends a general query to the switch, which forwards the query to ports 2 through 5, which are 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. The switch CPU uses the information in the IGMP report to set up a forwarding-table entry, as shown in Table 1-1, that includes the port numbers connected to Host 1 and the router.

Table 1-1 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 1-2), the CPU receives that message and adds the port number of Host 4 to the forwarding table as shown in Table 1-2. Note that because the forwarding table directs IGMP messages only to 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 1-2 Second Host Joining a Multicast Group




Table 1-2 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 a group-specific query to learn 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

Immediate Leave is only supported on IGMP Version 2 hosts.

The switch uses IGMP snooping Immediate Leave to remove from the forwarding table an interface that sends a leave message without the switch sending group-specific 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 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 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 inadvertently be dropped.

For configuration steps, see the [“Enabling IGMP Immediate Leave”](#) section on page 1-10.

IGMP Configurable-Leave Timer

You can configure the time that the switch waits after sending a group-specific query to determine if hosts are still interested in a specific multicast group. The IGMP leave response time can be configured from 100 to 5000 milliseconds. The timer can be set either globally or on a per-VLAN basis. The VLAN configuration of the leave time overrides the global configuration.

For configuration steps, see the [“Configuring the IGMP Leave Timer”](#) section on page 1-11.

IGMP Report Suppression

**Note**

IGMP report suppression is supported only when the multicast query has IGMPv1 and IGMPv2 reports. This feature is not supported when the query includes IGMPv3 reports.

The switch uses IGMP report suppression to forward only one IGMP report per multicast router query to multicast devices. When IGMP router suppression is enabled (the default), the switch sends the first IGMP report from all hosts for a group to all the multicast routers. The switch does not send the remaining IGMP reports for the group to the multicast routers. This feature prevents duplicate reports from being sent to the multicast devices.

If the multicast router query includes requests only for IGMPv1 and IGMPv2 reports, the switch forwards only the first IGMPv1 or IGMPv2 report from all hosts for a group to all the multicast routers.

If the multicast router query also includes requests for IGMPv3 reports, the switch forwards all IGMPv1, IGMPv2, and IGMPv3 reports for a group to the multicast devices.

If you disable IGMP report suppression, all IGMP reports are forwarded to the multicast routers. For configuration steps, see the [“Disabling IGMP Report Suppression”](#) section on page 1-15.

IGMP Snooping and Switch Stacks

IGMP snooping functions across the switch stack; that is, IGMP control information from one switch is distributed to all switches in the stack. (See [Chapter 1, “Managing Switch Stacks,”](#) for more information about switch stacks.) Regardless of the stack member through which IGMP multicast data enters the stack, the data reaches the hosts that have registered for that group.

If a switch in the stack fails or is removed from the stack, only the members of the multicast group that are on that switch will not receive the multicast data. All other members of a multicast group on other switches in the stack continue to receive multicast data streams. However, multicast groups that are common for both Layer 2 and Layer 3 (IP multicast routing) might take longer to converge if the stack master is removed.

Configuring IGMP Snooping

IGMP snooping allows switches to examine IGMP packets and make forwarding decisions based on their content. These sections contain this configuration information:

- [Default IGMP Snooping Configuration, page 1-7](#)
- [Enabling or Disabling IGMP Snooping, page 1-7](#)
- [Setting the Snooping Method, page 1-8](#)
- [Configuring a Multicast Router Port, page 1-9](#)
- [Configuring a Host Statically to Join a Group, page 1-10](#)
- [Enabling IGMP Immediate Leave, page 1-10](#)
- [Configuring the IGMP Leave Timer, page 1-11](#)
- [Configuring TCN-Related Commands, page 1-12](#)
- [Configuring the IGMP Snooping Querier, page 1-14](#)
- [Disabling IGMP Report Suppression, page 1-15](#)

Default IGMP Snooping Configuration

Table 1-3 shows the default IGMP snooping configuration.

Table 1-3 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
TCN ¹ flood query count	2
TCN query solicitation	Disabled
IGMP snooping querier	Disabled
IGMP report suppression	Enabled

1. TCN = Topology Change Notification

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 1001 and 1006 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 to 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.

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 1001 and 1006 to 4094. Specify the multicast router learning method: <ul style="list-style-type: none"> • cgmp—Listen for CGMP packets. This method is useful for reducing control traffic. • pim-dvmrp—Snoop on IGMP queries and PIM-DVMRP packets. This is the default.
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.

To return to the default learning method, use the **no ip igmp snooping vlan *vlan-id* mrouter learn cgmp** global configuration command.

This example shows how to configure IGMP snooping to use CGMP packets as the learning method:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 1 mrouter learn cgmp
Switch(config)# end
```

Configuring a Multicast Router Port

To add a multicast router port (add a static connection to a multicast router), use the **ip igmp snooping vlan mrouter** global configuration command on the switch.

Beginning in privileged EXEC mode, follow these steps to enable a static connection to a multicast router:

	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 the interface to the multicast router. <ul style="list-style-type: none"> • The VLAN ID range is 1 to 1001 and 1006 to 4094. • The interface can be a physical interface or a port channel. The port-channel range is 1 to 6.
Step 3	end	Return to privileged EXEC mode.

	Command	Purpose
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:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 200 mrouter interface gigabitethernet0/2
Switch(config)# end
```

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: <ul style="list-style-type: none"> • <i>vlan-id</i> is the multicast group VLAN ID. The range is 1 to 1001 and 1006 to 4094. • <i>ip_address</i> is the group IP address. • <i>interface-id</i> is the member port. It can be a physical interface or a port channel (1 to 6).
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip igmp snooping groups	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 a port:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 105 static 224.2.4.12 interface gigabitethernet0/1
Switch(config)# end
```

Enabling IGMP Immediate Leave

When you enable IGMP Immediate Leave, the switch immediately removes a port when it detects an IGMP Version 2 leave message on that port. You should only use the Immediate-Leave feature when there is a single receiver present on every port in the VLAN.



Note Immediate Leave is supported only on IGMP Version 2 hosts.

Beginning in privileged EXEC mode, follow these steps to enable IGMP Immediate Leave:

	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 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 interface.
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 on VLAN 130:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 130 immediate-leave
Switch(config)# end
```

Configuring the IGMP Leave Timer

Follows these guidelines when configuring the IGMP leave timer:

- You can configure the leave time globally or on a per-VLAN basis.
- Configuring the leave time on a VLAN overrides the global setting.
- The default leave time is 1000 milliseconds.
- The IGMP configurable leave time is only supported on hosts running IGMP Version 2.
- The actual leave latency in the network is usually the configured leave time. However, the leave time *might* vary around the configured time, depending on real-time CPU load conditions, network delays and the amount of traffic sent through the interface.

Beginning in privileged EXEC mode, follow these steps to enable the IGMP configurable-leave timer:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip igmp snooping last-member-query-interval <i>time</i>	Configure the IGMP leave timer globally. The range is 100 to 32768 milliseconds. The default is 1000 seconds.
Step 3	ip igmp snooping vlan <i>vlan-id</i> last-member-query-interval <i>time</i>	(Optional) Configure the IGMP leave time on the VLAN interface. The range is 100 to 32768 milliseconds. Note Configuring the leave time on a VLAN overrides the globally configured timer.
Step 4	end	Return to privileged EXEC mode.

	Command	Purpose
Step 5	show ip igmp snooping	(Optional) Display the configured IGMP leave time.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To globally reset the IGMP leave timer to the default setting, use the **no ip igmp snooping last-member-query-interval** global configuration command.

To remove the configured IGMP leave-time setting from the specified VLAN, use the **no ip igmp snooping vlan *vlan-id* last-member-query-interval** global configuration command.

Configuring TCN-Related Commands

These sections describe how to control flooded multicast traffic during a TCN event:

- [Controlling the Multicast Flooding Time After a TCN Event, page 1-12](#)
- [Recovering from Flood Mode, page 1-13](#)
- [Disabling Multicast Flooding During a TCN Event, page 1-13](#)

Controlling the Multicast Flooding Time After a TCN Event

You can control the time that multicast traffic is flooded after a TCN event by using the **ip igmp snooping tcn flood query count** global configuration command. This command configures the number of general queries for which multicast data traffic is flooded after a TCN event. Some examples of TCN events are when the client changed its location and the receiver is on same port that was blocked but is now forwarding, and when a port went down without sending a leave message.

If you set the TCN flood query count to 1 by using the **ip igmp snooping tcn flood query count** command, the flooding stops after receiving 1 general query. If you set the count to 7, the flooding until 7 general queries are received. Groups are relearned based on the general queries received during the TCN event.

Beginning in privileged EXEC mode, follow these steps to configure the TCN flood query count:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip igmp snooping tcn flood query count <i>count</i>	Specify the number of IGMP general queries for which the multicast traffic is flooded. The range is 1 to 10. By default, the flooding query count is 2.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip igmp snooping	Verify the TCN settings.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default flooding query count, use the **no ip igmp snooping tcn flood query count** global configuration command.

Recovering from Flood Mode

When a topology change occurs, the spanning-tree root sends a special IGMP leave message (also known as global leave) with the group multicast address 0.0.0.0. However, when you enable the **ip igmp snooping tcn query solicit** global configuration command, the switch sends the global leave message whether or not it is the spanning-tree root. When the router receives this special leave, it immediately sends general queries, which expedite the process of recovering from the flood mode during the TCN event. Leaves are always sent if the switch is the spanning-tree root regardless of this configuration command. By default, query solicitation is disabled.

Beginning in privileged EXEC mode, follow these steps to enable the switch to send the global leave message whether or not it is the spanning-tree root:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip igmp snooping tcn query solicit	Send an IGMP leave message (global leave) to speed the process of recovering from the flood mode caused during a TCN event. By default, query solicitation is disabled.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip igmp snooping	Verify the TCN settings.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default query solicitation, use the **no ip igmp snooping tcn query solicit** global configuration command.

Disabling Multicast Flooding During a TCN Event

When the switch receives a TCN, multicast traffic is flooded to all the ports until 2 general queries are received. If the switch has many ports with attached hosts that are subscribed to different multicast groups, this flooding might exceed the capacity of the link and cause packet loss. You can use the **ip igmp snooping tcn flood** interface configuration command to control this behavior.

Beginning in privileged EXEC mode, follow these steps to disable multicast flooding on an interface:

	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	no ip igmp snooping tcn flood	Disable the flooding of multicast traffic during a spanning-tree TCN event. By default, multicast flooding is enabled on an interface.
Step 4	exit	Return to privileged EXEC mode.
Step 5	show ip igmp snooping	Verify the TCN settings.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To re-enable multicast flooding on an interface, use the **ip igmp snooping tcn flood** interface configuration command.

Configuring the IGMP Snooping Querier

Follow these guidelines when configuring the IGMP snooping querier:

- Configure the VLAN in global configuration mode.
- Configure an IP address on the VLAN interface. When enabled, the IGMP snooping querier uses the IP address as the query source address.
- If there is no IP address configured on the VLAN interface, the IGMP snooping querier tries to use the configured global IP address for the IGMP querier. If there is no global IP address specified, the IGMP querier tries to use the VLAN switch virtual interface (SVI) IP address (if one exists). If there is no SVI IP address, the switch uses the first available IP address configured on the switch. The first IP address available appears in the output of the **show ip interface** privileged EXEC command. The IGMP snooping querier does not generate an IGMP general query if it cannot find an available IP address on the switch.
- The IGMP snooping querier supports IGMP Versions 1 and 2.
- When administratively enabled, the IGMP snooping querier moves to the nonquerier state if it detects the presence of a multicast router in the network.
- When it is administratively enabled, the IGMP snooping querier moves to the operationally disabled state under these conditions:
 - IGMP snooping is disabled in the VLAN.
 - PIM is enabled on the SVI of the corresponding VLAN.

Beginning in privileged EXEC mode, follow these steps to enable the IGMP snooping querier feature in a VLAN:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip igmp snooping querier	Enable the IGMP snooping querier.
Step 3	ip igmp snooping querier address <i>ip_address</i>	(Optional) Specify an IP address for the IGMP snooping querier. If you do not specify an IP address, the querier tries to use the global IP address configured for the IGMP querier. Note The IGMP snooping querier does not generate an IGMP general query if it cannot find an IP address on the switch.
Step 4	ip igmp snooping querier query-interval <i>interval-count</i>	(Optional) Set the interval between IGMP queries. The range is 1 to 18000 seconds.
Step 5	ip igmp snooping querier tcn query [count <i>count</i> interval <i>interval</i>]	(Optional) Set the time between Topology Change Notification (TCN) queries. The count range is 1 to 10. The interval range is 1 to 255 seconds.
Step 6	ip igmp snooping querier timer expiry <i>timeout</i>	(Optional) Set the length of time until the IGMP querier expires. The range is 60 to 300 seconds.
Step 7	ip igmp snooping querier version <i>version</i>	(Optional) Select the IGMP version number that the querier feature uses. Select 1 or 2.
Step 8	end	Return to privileged EXEC mode.
Step 9	show ip igmp snooping vlan <i>vlan-id</i>	(Optional) Verify that the IGMP snooping querier is enabled on the VLAN interface. The VLAN ID range is 1 to 1001 and 1006 to 4094.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to set the IGMP snooping querier source address to 10.0.0.64:

```
Switch# configure terminal
Switch(config)# ip igmp snooping querier 10.0.0.64
Switch(config)# end
```

This example shows how to set the IGMP snooping querier maximum response time to 25 seconds:

```
Switch# configure terminal
Switch(config)# ip igmp snooping querier query-interval 25
Switch(config)# end
```

This example shows how to set the IGMP snooping querier timeout to 60 seconds:

```
Switch# configure terminal
Switch(config)# ip igmp snooping querier timeout expiry 60
Switch(config)# end
```

This example shows how to set the IGMP snooping querier feature to version 2:

```
Switch# configure terminal
Switch(config)# no ip igmp snooping querier version 2
Switch(config)# end
```

Disabling IGMP Report Suppression



Note

IGMP report suppression is supported only when the multicast query has IGMPv1 and IGMPv2 reports. This feature is not supported when the query includes IGMPv3 reports.

IGMP report suppression is enabled by default. When it is enabled, the switch forwards only one IGMP report per multicast router query. When report suppression is disabled, all IGMP reports are forwarded to the multicast routers.

Beginning in privileged EXEC mode, follow these steps to disable IGMP report suppression:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no ip igmp snooping report-suppression	Disable IGMP report suppression.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip igmp snooping	Verify that IGMP report suppression is disabled.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To re-enable IGMP report suppression, use the **ip igmp snooping report-suppression** global configuration command.

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 1-4.

Table 1-4 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. The VLAN ID range is 1 to 1001 and 1006 to 4094.
show ip igmp snooping groups [count dynamic [count] user [count]]	Display multicast table information for the switch or about a specific parameter: <ul style="list-style-type: none"> • count—Display the total number of entries for the specified command options instead of the actual entries. • dynamic—Display entries learned through IGMP snooping. • user—Display only the user-configured multicast entries.
show ip igmp snooping groups vlan <i>vlan-id</i> [<i>ip_address</i> count dynamic [count] user [count]]	Display multicast table information for a multicast VLAN or about a specific parameter for the VLAN: <ul style="list-style-type: none"> • <i>vlan-id</i>—The VLAN ID range is 1 to 1001 and 1006 to 4094. • count—Display the total number of entries for the specified command options instead of the actual entries. • dynamic—Display entries learned through IGMP snooping. • <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.
show ip igmp snooping querier [vlan <i>vlan-id</i>]	Display information about the IP address and receiving port for the most-recently received IGMP query messages in the VLAN. (Optional) Enter vlan <i>vlan-id</i> to display information for a single VLAN.
show ip igmp snooping querier [vlan <i>vlan-id</i>] detail	Display information about the IP address and receiving port of the most-recently received IGMP query message in the VLAN and the configuration and operational state of the IGMP snooping querier in the VLAN.

For more information about the keywords and options in these commands, see the command reference for this release.

Understanding Multicast VLAN Registration

**Note**

To use MVR, the switch must be running the LAN Base image.

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 that MVR hosts have joined, either by IGMP reports or by MVR static configuration. IGMP reports received from MVR hosts are never forwarded from 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 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 stack is supported.


Receiver ports and source ports can be on different switches in a switch stack. Multicast data sent on the multicast VLAN is forwarded to all MVR receiver ports across the stack. When a new switch is added to a stack, by default it has no receiver ports.

If a switch fails or is removed from the stack, only those receiver ports belonging to that switch will not receive the multicast data. All other receiver ports on other switches continue to receive the multicast data.

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 1-3 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 1-3 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.

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. Switch B. 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

- [Default MVR Configuration, page 1-19](#)
- [MVR Configuration Guidelines and Limitations, page 1-20](#)
- [Configuring MVR Global Parameters, page 1-20](#)
- [Configuring MVR Interfaces, page 1-21](#)

Default MVR Configuration

Table 1-5 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
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 can coexist with IGMP snooping on a switch.
- MVR data received on an MVR receiver port is not forwarded to MVR source ports.
- MVR does not support IGMPv3 messages.

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, see 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.
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 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 1001 and 1006 to 4094. The default is VLAN 1.

	Command	Purpose
Step 6	mvr mode {dynamic compatible}	(Optional) Specify the MVR mode of operation: <ul style="list-style-type: none"> • dynamic—Allows dynamic MVR membership on source ports. • compatible—Is compatible with Catalyst 3500 XL and Catalyst 2900 XL switches and does not support IGMP dynamic joins on source ports. The default is compatible mode.
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 interface-id	Specify the Layer 2 port to configure, and enter interface configuration mode.
Step 4	mvr type {source receiver}	Configure an MVR port as one of these: <ul style="list-style-type: none"> • source—Configure uplink ports that receive and send multicast data as source ports. Subscribers cannot be directly connected to source ports. All source ports on a switch belong to the single multicast VLAN. • receiver—Configure a port as a receiver port if it is a subscriber port and should only receive multicast data. It does not receive data unless it becomes a member of the multicast group, either statically or by using IGMP leave and join messages. Receiver ports cannot belong to the multicast VLAN. The default configuration is as a non-MVR port. If you attempt to configure a non-MVR port with MVR characteristics, the operation fails.

	Command	Purpose
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 show mvr interface 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 interface to its default settings, use the **no mvr [type | immediate | vlan *vlan-id* | group]** interface configuration commands.

This example shows how to configure a port as a receiver port, statically configure the port to receive multicast traffic sent to the multicast group address, configure Immediate Leave on the port, 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
----      -
Gi0/2    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 1-6](#) to display MVR configuration:

Table 1-6 Commands for Displaying MVR Information

Command	Purpose
<code>show mvr</code>	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.
<code>show mvr interface [interface-id] [members [vlan vlan-id]]</code>	<p>Displays all MVR interfaces and their MVR configurations.</p> <p>When a specific interface is entered, displays this information:</p> <ul style="list-style-type: none"> • Type—Receiver or Source • Status—One of these: <ul style="list-style-type: none"> – Active means the port is part of a VLAN. – Up/Down means that the port is forwarding or nonforwarding. – Inactive means that the port is not part of any VLAN. • Immediate Leave—Enabled or Disabled <p>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 1001 and 1006 to 4094.</p>
<code>show mvr members [ip-address]</code>	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 and Throttling

In some environments, for example, metropolitan or multiple-dwelling unit (MDU) installations, you might want to control the set of multicast groups to which a user on a switch port can belong. You can control the distribution of multicast services, such as IP/TV, based on some type of subscription or service plan. You might also want to limit the number of multicast groups to which a user on a switch port can belong.

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. You can also set the maximum number of IGMP groups that a Layer 2 interface can join.

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 applicable only to the dynamic learning of IP multicast group addresses, not static configuration.

With the IGMP throttling feature, you can set the maximum number of IGMP groups that a Layer 2 interface can join. If the maximum number of IGMP groups is set, the IGMP snooping forwarding table contains the maximum number of entries, and the interface receives an IGMP join report, you can configure an interface to drop the IGMP report or to replace the randomly selected multicast entry with the received IGMP report.



Note

IGMPv3 join and leave messages are not supported on switches running IGMP filtering.

These sections contain this configuration information:

- [Default IGMP Filtering and Throttling Configuration, page 1-24](#)
- [Configuring IGMP Profiles, page 1-24](#) (optional)
- [Applying IGMP Profiles, page 1-25](#) (optional)
- [Setting the Maximum Number of IGMP Groups, page 1-26](#) (optional)
- [Configuring the IGMP Throttling Action, page 1-27](#) (optional)

Default IGMP Filtering and Throttling Configuration

[Table 1-7](#) shows the default IGMP filtering configuration.

Table 1-7 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

When the maximum number of groups is in forwarding table, the default IGMP throttling action is to deny the IGMP report. For configuration guidelines, see the [“Configuring the IGMP Throttling Action” section on page 1-27](#).

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.
- **exit**: Exits from igmp-profile configuration mode.
- **no**: Negates a command or returns to 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>	Assign a number to the profile you are configuring, and enter IGMP profile configuration mode. The profile number range is 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 only to Layer 2 access ports. 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 have only 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>	Specify the physical interface, and enter interface configuration mode. 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 range is 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 gigabitethernet0/2
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 max-groups** interface configuration command. Use the **no** form of this command to set the maximum back to the default, which is 208.

You can use this command on a logical EtherChannel interface but cannot use it on ports that belong to an EtherChannel port group.

Beginning in privileged EXEC mode, follow these steps to set the maximum number of IGMP groups in the forwarding table:

	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. The interface can be a Layer 2 port that does not belong to an EtherChannel group or a EtherChannel interface.
Step 3	ip igmp max-groups <i>number</i>	Set the maximum number of IGMP groups that the interface can join. The range is 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 to 25 the number of IGMP groups that a port can join.

```
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# ip igmp max-groups 25
Switch(config-if)# end
```

Configuring the IGMP Throttling Action

After you set the maximum number of IGMP groups that a Layer 2 interface can join, you can configure an interface to replace the existing group with the new group for which the IGMP report was received by using the **ip igmp max-groups action replace** interface configuration command. Use the **no** form of this command to return to the default, which is to drop the IGMP join report.

Follow these guidelines when configuring the IGMP throttling action:

- You can use this command on a logical EtherChannel interface but cannot use it on ports that belong to an EtherChannel port group.
- When the maximum group limitation is set to the default (no maximum), entering the **ip igmp max-groups action {deny | replace}** command has no effect.
- If you configure the throttling action and set the maximum group limitation after an interface has added multicast entries to the forwarding table, the forwarding-table entries are either aged out or removed, depending on the throttling action.
 - If you configure the throttling action as **deny**, the entries that were previously in the forwarding table are not removed but are aged out. After these entries are aged out and the maximum number of entries is in the forwarding table, the switch drops the next IGMP report received on the interface.
 - If you configure the throttling action as **replace**, the entries that were previously in the forwarding table are removed. When the maximum number of entries is in the forwarding table, the switch replaces a randomly selected entry with the received IGMP report.

To prevent the switch from removing the forwarding-table entries, you can configure the IGMP throttling action before an interface adds entries to the forwarding table.

Beginning in privileged EXEC mode, follow these steps to configure the throttling action when the maximum number of entries is in the forwarding table:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the physical interface to be configured, and enter interface configuration mode. The interface can be a Layer 2 port that does not belong to an EtherChannel group or an EtherChannel interface. The interface cannot be a trunk port.
Step 3	ip igmp max-groups action {deny replace}	When an interface receives an IGMP report and the maximum number of entries is in the forwarding table, specify the action that the interface takes: <ul style="list-style-type: none"> • deny—Drop the report. • replace—Replace the existing group with the new group for which the IGMP report was received.

	Command	Purpose
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 return to the default action of dropping the report, use the **no ip igmp max-groups action** interface configuration command.

Displaying IGMP Filtering and Throttling 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. You can also display the IGMP throttling configuration for all interfaces on the switch or for a specified interface.

Use the privileged EXEC commands in [Table 1-8](#) to display IGMP filtering and throttling configuration:

Table 1-8 Commands for Displaying IGMP Filtering and Throttling Configuration

Command	Purpose
show ip igmp profile [<i>profile number</i>]	Displays the specified IGMP profile or all the IGMP profiles defined on the switch.
show running-config [interface <i>interface-id</i>]	Displays the configuration of the specified interface or the configuration of 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.



CHAPTER 1

Configuring Port-Based Traffic Control

This chapter describes how to configure the port-based traffic control features on the Catalyst 2960 and 2960-S switches. Unless otherwise noted, the term *switch* refers to a standalone switch and a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release.

This chapter consists of these sections:

- [Configuring Storm Control, page 1-1](#)
- [Configuring Protected Ports, page 1-6](#)
- [Configuring Port Blocking, page 1-7](#)
- [Configuring Port Security, page 1-8](#)
- [Configuring Protocol Storm Protection, page 1-18](#)
- [Displaying Port-Based Traffic Control Settings, page 1-19](#)

Configuring Storm Control

- [Understanding Storm Control, page 1-1](#)
- [Default Storm Control Configuration, page 1-3](#)
- [Configuring Storm Control and Threshold Levels, page 1-3](#)
- [Configuring Small-Frame Arrival Rate, page 1-5](#)

Understanding Storm Control

Storm control prevents traffic 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, mistakes in network configurations, or users issuing a denial-of-service attack can cause a storm.

Storm control (or traffic suppression) monitors packets passing from an interface to the switching bus and determines if the packet is unicast, multicast, or broadcast. The switch counts the number of packets of a specified type received within the 1-second time interval and compares the measurement with a predefined suppression-level threshold.

Storm control uses one of these methods to measure traffic activity:

- Bandwidth as a percentage of the total available bandwidth of the port that can be used by the broadcast, multicast, or unicast traffic
- Traffic rate in packets per second at which broadcast, multicast, or unicast packets are received.
- Traffic rate in bits per second at which broadcast, multicast, or unicast packets are received.
- Traffic rate in packets per second and for small frames. This feature is enabled globally. The threshold for small frames is configured for each interface.

With each method, the port blocks traffic when the rising threshold is reached. The port remains blocked until the traffic rate drops below the falling threshold (if one is specified) and then resumes normal forwarding. If the falling suppression level is not specified, the switch blocks all traffic until the traffic rate drops below the rising suppression level. In general, the higher the level, the less effective the protection against broadcast storms.




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.

The graph in [Figure 1-1](#) 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 1-1 Broadcast Storm Control Example



The combination of the storm-control suppression level and the 1-second time interval controls 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 1-second time interval during which traffic activity is measured can affect the behavior of storm control.

You use the **storm-control** interface configuration commands to set the threshold value for each traffic type.

Default Storm Control Configuration

By default, unicast, broadcast, and multicast storm control are disabled on the switch interfaces; that is, the suppression level is 100 percent.

Configuring Storm Control and Threshold Levels

You configure storm control on a port and enter the threshold level that you want to be used for a particular type of 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 on physical interfaces. You can also configure storm control on an EtherChannel. When storm control is configured on an EtherChannel, the storm control settings propagate to the EtherChannel physical interfaces.

Beginning in privileged EXEC mode, follow these steps to storm control and threshold levels:

	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	storm-control { broadcast multicast unicast } level { <i>level</i> [<i>level-low</i>] bps <i>bps</i> [<i>bps-low</i>] pps <i>pps</i> [<i>pps-low</i>]}	<p>Configure broadcast, multicast, or unicast storm control. By default, storm control is disabled.</p> <p>The keywords have these meanings:</p> <ul style="list-style-type: none"> For <i>level</i>, specify the rising threshold level for broadcast, multicast, or unicast traffic as a percentage (up to two decimal places) of the bandwidth. The port blocks traffic when the rising threshold is reached. The range is 0.00 to 100.00. (Optional) For <i>level-low</i>, specify the falling threshold level as a percentage (up to two decimal places) of the bandwidth. This value must be less than or equal to the rising suppression value. The port forwards traffic when traffic drops below this level. If you do not configure a falling suppression level, it is set to the rising suppression level. The range is 0.00 to 100.00. <p>If you set the threshold to the maximum value (100 percent), no limit is placed on the traffic. If you set the threshold to 0.0, all broadcast, multicast, and unicast traffic on that port is blocked.</p> <ul style="list-style-type: none"> For bps <i>bps</i>, specify the rising threshold level for broadcast, multicast, or unicast traffic in bits per second (up to one decimal place). The port blocks traffic when the rising threshold is reached. The range is 0.0 to 10000000000.0. (Optional) For <i>bps-low</i>, specify the falling threshold level in bits per second (up to one decimal place). It can be less than or equal to the rising threshold level. The port forwards traffic when traffic drops below this level. The range is 0.0 to 10000000000.0. For pps <i>pps</i>, specify the rising threshold level for broadcast, multicast, or unicast traffic in packets per second (up to one decimal place). The port blocks traffic when the rising threshold is reached. The range is 0.0 to 10000000000.0. (Optional) For <i>pps-low</i>, specify the falling threshold level in packets per second (up to one decimal place). It can be less than or equal to the rising threshold level. The port forwards traffic when traffic drops below this level. The range is 0.0 to 10000000000.0. <p>For BPS and PPS settings, you can use metric suffixes such as k, m, and g for large number thresholds.</p>
Step 4	storm-control action { shutdown trap }	<p>Specify the action to be taken when a storm is detected. The default is to filter out the traffic and not to send traps.</p> <ul style="list-style-type: none"> Select the shutdown keyword to error-disable the port during a storm. Select the trap keyword to generate an SNMP trap when a storm is detected.
Step 5	end	Return to privileged EXEC mode.
Step 6	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 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable storm control, use the **no storm-control {broadcast | multicast | unicast} level** interface configuration command.

This example shows how to enable unicast storm control on a port with an 87-percent rising suppression level and a 65-percent falling suppression level:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# storm-control unicast level 87 65
```

This example shows how to enable broadcast address storm control on a port to a level of 20 percent. When the broadcast traffic exceeds the configured level of 20 percent of the total available bandwidth of the port within the traffic-storm-control interval, the switch drops all broadcast traffic until the end of the traffic-storm-control interval:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# storm-control broadcast level 20
```

Configuring Small-Frame Arrival Rate

Incoming VLAN-tagged packets smaller than 67 bytes are considered *small frames*. They are forwarded by the switch, but they do not cause the switch storm-control counters to increment. In Cisco IOS Release 12.2(44)SE and later, you can configure a port to be error disabled if small frames arrive at a specified rate (threshold).

You globally enable the small-frame arrival feature on the switch and then configure the small-frame threshold for packets on each interface. Packets smaller than the minimum size and arriving at a specified rate (the threshold) are dropped since the port is error disabled.

If the **errdisable recovery cause small-frame** global configuration command is entered, the port is re-enabled after a specified time. (You specify the recovery time by using **errdisable recovery** global configuration command.)



Note

Small Frame Arrival-Rate is supported by 2960 only.

Beginning in privileged EXEC mode, follow these steps to configure the threshold level for each interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	errdisable detect cause small-frame	Enable the small-frame rate-arrival feature on the switch.
Step 3	errdisable recovery interval <i>interval</i>	(Optional) Specify the time to recover from the specified error-disabled state.
Step 4	errdisable recovery cause small-frame	(Optional) Configure the recovery time for error-disabled ports to be automatically re-enabled after they are error disabled by the arrival of small frames
Step 5	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be configured.

	Command	Purpose
Step 6	<code>small violation-rate pps</code>	Configure the threshold rate for the interface to drop incoming packets and error disable the port. The range is 1 to 10,000 packets per second (pps)
Step 7	<code>end</code>	Return to privileged EXEC mode.
Step 8	<code>show interfaces interface-id</code>	Verify the configuration.
Step 9	<code>copy running-config startup-config</code>	(Optional) Save your entries in the configuration file.

This example shows how to enable the small-frame arrival-rate feature, configure the port recovery time, and configure the threshold for error disabling a port:

```
Switch# configure terminal
Switch# errdisable detect cause small-frame
Switch# errdisable recovery cause small-frame
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# small-frame violation rate 10000
Switch(config-if)# end
```

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. Data traffic cannot be forwarded between protected ports at Layer 2; only control traffic, such as PIM packets, is forwarded because these packets are processed by the CPU and forwarded in software. All data 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.

Because a switch stack represents a single logical switch, Layer 2 traffic is not forwarded between any protected ports in the switch stack, whether they are on the same or different switches in the stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

These sections contain this configuration information:

- [Default Protected Port Configuration, page 1-6](#)
- [Protected Port Configuration Guidelines, page 1-7](#)
- [Configuring a Protected Port, page 1-7](#)

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 Ethernet port 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>	Specify the interface to be configured, and enter interface configuration mode.
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 a port as a protected port:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
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.



Note

With multicast traffic, the port blocking feature blocks only pure Layer 2 packets. Multicast packets that contain IPv4 or IPv6 information in the header are not blocked.

- [Default Port Blocking Configuration, page 1-8](#)
- [Blocking Flooded Traffic on an Interface, page 1-8](#)

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 or an EtherChannel group. 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 unicast packets and Layer 2 multicast packets out of an interface:

	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	switchport block multicast	Block unknown multicast forwarding out of the port. Note Only pure Layer 2 multicast traffic is blocked. Multicast packets that contain IPv4 or IPv6 information in the header are not blocked.
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 Layer 2 multicast flooding on a port:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
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 contain this conceptual and configuration information:

- [Understanding Port Security, page 1-9](#)
- [Default Port Security Configuration, page 1-11](#)
- [Port Security Configuration Guidelines, page 1-11](#)
- [Enabling and Configuring Port Security, page 1-13](#)
- [Enabling and Configuring Port Security Aging, page 1-16](#)
- [Port Security and Switch Stacks, page 1-18](#)

Understanding Port Security

- [Secure MAC Addresses, page 1-9](#)
- [Security Violations, page 1-10](#)

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 set by the maximum number of available MAC addresses allowed in the system. This number is 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 four 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. When a secure port is in the error-disabled state, you can bring it out of this state by entering the **errdisable recovery cause psecure-violation** global configuration command, or you can manually re-enable it by entering the **shutdown** and **no shutdown** interface configuration commands. This is the default mode.
- **shutdown vlan**—Use to set the security violation mode per-VLAN. In this mode, the VLAN is error disabled instead of the entire port when a violation occurs

Table 1-1 shows the violation mode and the actions taken when you configure an interface for port security.

Table 1-1 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

Table 1-1 Security Violation Mode Actions (continued)

Violation Mode	Traffic is forwarded ¹	Sends SNMP trap	Sends syslog message	Displays error message ²	Violation counter increments	Shuts down port
shutdown	No	No	No	No	Yes	Yes
shutdown vlan	No	No	Yes	No	Yes	No ³

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.
3. Shuts down only the VLAN on which the violation occurred.

Default Port Security Configuration

Table 1-2 shows the default port security configuration for an interface.

Table 1-2 Default Port Security Configuration

Feature	Default Setting
Port security	Disabled on a port.
Sticky address learning	Disabled.
Maximum number of 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.

Port Security 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.



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, set the maximum allowed secure addresses on the port to two. When the port is connected to a Cisco IP phone, the IP phone requires one MAC address. The Cisco IP phone address is learned on the voice

VLAN, but is not learned on the access VLAN. If you connect a single PC to the Cisco IP phone, no additional MAC addresses are required. If you connect more than one PC to the Cisco IP phone, you must configure enough secure addresses to allow one for each PC and one for the phone.

- When a trunk port configured with port security and assigned to an access VLAN for data traffic and to a voice VLAN for voice traffic, entering the **switchport voice** and **switchport priority extend** interface configuration commands has no effect.

When a connected device uses the same MAC address to request an IP address for the access VLAN and then an IP address for the voice VLAN, only the access VLAN is assigned an IP address.

- When configuring port security, first specify the total number of MAC addresses you want to allow, by using the **switchport port-security maximum** interface configuration command and then configure the number of access VLANs (**switchport port-security vlan access** interface configuration command) and voice VLANs (**switchport port-security vlan voice** interface configuration command) you want to allow. If you do not specify the total number first, the system returns to the default setting (1 MAC address).
- 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.

Table 1-3 summarizes port security compatibility with other port-based features.

Table 1-3 Port Security Compatibility with Other Switch Features

Type of Port or Feature on Port	Compatible with Port Security
DTP ¹ port ²	No
Trunk port	Yes
Dynamic-access port ³	No
SPAN source port	Yes
SPAN destination port	No
EtherChannel	No
Protected port	Yes
IEEE 802.1x port	Yes
Voice VLAN port ⁴	Yes
Flex Links	Yes

1. DTP = Dynamic Trunking Protocol
2. A port configured with the **switchport mode dynamic** interface configuration command.
3. A VLAN Query Protocol (VQP) port configured with the **switchport access vlan dynamic** interface configuration command.
4. 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.

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>	Specify the interface to be configured, and enter interface configuration mode.
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 voice vlan <i>vlan-id</i>	Enable voice VLAN on a port. <i>vlan-id</i> —Specify the VLAN to be used for voice traffic.
Step 5	switchport port-security	Enable port security on the interface.
Step 6	switchport port-security [maximum <i>value</i> [vlan { <i>vlan-list</i> { access voice }}]]	(Optional) maximum — Specify the maximum number of secure MAC addresses on the port. By default only 1 MAC address is allowed. The maximum number of secure MAC addresses that you can configure on a switch is set by the maximum number of available MAC addresses allowed in the system. This number is the total of available MAC addresses, including those used for other Layer 2 functions and any other secure MAC addresses configured on interfaces. (Optional) vlan —set a per-VLAN maximum value Enter one of these options after you enter the vlan keyword: <ul style="list-style-type: none"> <i>vlan-list</i>—On a trunk port, you can set a per-VLAN maximum value on a range of VLANs separated by a hyphen or a series of VLANs separated by commas. For nonspecified VLANs, the per-VLAN maximum value is used. access—On an access port, specify the VLAN as an access VLAN. voice—On an access port, specify the VLAN as a voice VLAN. Note The voice keyword is available only if a voice VLAN is configured on a port and if that port is not the access VLAN. If an interface is configured for voice VLAN, configure a maximum of two secure MAC addresses.

Command	Purpose
<p>Step 7</p> <pre>switchport port-security [violation {protect restrict shutdown shutdown vlan}]</pre>	<p>(Optional) Set the violation mode, the action to be taken when a security violation is detected, as one of these:</p> <ul style="list-style-type: none"> • protect—When the number of 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. <p>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.</p> <ul style="list-style-type: none"> • 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—When a violation occurs, the interface is error disabled, the port LED turns off, and the violation counter increments. • shutdown vlan—Use to set the security violation mode per VLAN. In this mode, the VLAN is error disabled instead of the entire port when a violation occurs. <p>Note When a secure port is in the error-disabled state, you can bring it out of this state by entering the errdisable recovery cause psecure-violation global configuration command. You can manually re-enable it by entering the shutdown and no shutdown interface configuration commands or by using the clear errdisable interface vlan privileged EXEC command.</p>
<p>Step 8</p> <pre>switchport port-security [mac-address mac-address [vlan {vlan-id {access voice}}]</pre>	<p>(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.</p> <p>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.</p> <p>(Optional) vlan—set a per-VLAN maximum value.</p> <p>Enter one of these options after you enter the vlan keyword:</p> <ul style="list-style-type: none"> • vlan-id—On a trunk port, you can specify the VLAN ID and the MAC address. If you do not specify a VLAN ID, the native VLAN is used. • access—On an access port, specify the VLAN as an access VLAN. • voice—On an access port, specify the VLAN as a voice VLAN. <p>Note The voice keyword is available only if a voice VLAN is configured on a port and if that port is not the access VLAN. If an interface is configured for voice VLAN, configure a maximum of two secure MAC addresses.</p>

	Command	Purpose
Step 9	switchport port-security mac-address sticky	(Optional) Enable sticky learning on the interface.
Step 10	switchport port-security mac-address sticky [<i>mac-address</i> vlan { <i>vlan-id</i> { access voice }}]	<p>(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.</p> <p>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.</p> <p>(Optional) vlan—set a per-VLAN maximum value.</p> <p>Enter one of these options after you enter the vlan keyword:</p> <ul style="list-style-type: none"> <i>vlan-id</i>—On a trunk port, you can specify the VLAN ID and the MAC address. If you do not specify a VLAN ID, the native VLAN is used. access—On an access port, specify the VLAN as an access VLAN. voice—On an access port, specify the VLAN as a voice VLAN. <p>Note The voice keyword is available only if a voice VLAN is configured on a port and if that port is not the access VLAN.</p>
Step 11	end	Return to privileged EXEC mode.
Step 12	show port-security	Verify your entries.
Step 13	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.

Use the **clear port-security {all | configured | dynamic | sticky}** privileged EXEC command to delete from the MAC address table all secure addresses or all secure addresses of a specific type (configured, dynamic, or sticky) on the switch or on an interface.

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 a port 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 gigabitethernet0/1
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 a port:

```
Switch(config)# interface gigabitethernet0/2
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
```

This example shows how to enable sticky port security on a port, to manually configure MAC addresses for data VLAN and voice VLAN, and to set the total maximum number of secure addresses to 20 (10 for data VLAN and 10 for voice VLAN).

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# switchport access vlan 21
Switch(config-if)# switchport mode access
Switch(config-if)# switchport voice vlan 22
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security maximum 20
Switch(config-if)# switchport port-security violation restrict
Switch(config-if)# switchport port-security mac-address sticky
Switch(config-if)# switchport port-security mac-address sticky 0000.0000.0002
Switch(config-if)# switchport port-security mac-address 0000.0000.0003
Switch(config-if)# switchport port-security mac-address sticky 0000.0000.0001 vlan voice
Switch(config-if)# switchport port-security mac-address 0000.0000.0004 vlan voice
Switch(config-if)# switchport port-security maximum 10 vlan access
Switch(config-if)# switchport port-security maximum 10 vlan voice
```

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>	Specify the interface to be configured, and enter interface configuration mode.
Step 3	switchport port-security aging { static time <i>time</i> type { absolute inactivity }}	<p>Enable or disable static aging for the secure port, or set the aging time or type.</p> <p>Note The switch does not support port security aging of sticky secure addresses.</p> <p>Enter static to enable aging for statically configured secure addresses on this port.</p> <p>For <i>time</i>, specify the aging time for this port. The valid range is from 0 to 1440 minutes.</p> <p>For type, select one of these keywords:</p> <ul style="list-style-type: none"> • 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 a port:

```
Switch(config)# interface gigabitethernet0/1
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.

Port Security and Switch Stacks

When a switch joins a stack, the new switch receives the configured secure addresses. The new stack member downloads all dynamic secure addresses from the other stack members.

When a switch (either the stack master or a stack member) leaves the stack, the remaining stack members are notified, and the secure MAC addresses configured or learned by that switch are deleted from the secure MAC address table. For more information about switch stacks, see [Chapter 1, “Managing Switch Stacks.”](#)

```
Switch(config)# interface gigabitethernet0/1
```

Configuring Protocol Storm Protection

- [Understanding Protocol Storm Protection, page 1-18](#)
- [Default Protocol Storm Protection Configuration, page 1-18](#)
- [Enabling Protocol Storm Protection, page 1-19](#)

Understanding Protocol Storm Protection

When a switch is flooded with Address Resolution Protocol (ARP) or control packets, high CPU utilization can cause the CPU to overload. These issues can occur:

- Routing protocol can flap because the protocol control packets are not received, and neighboring adjacencies are dropped.
- Spanning Tree Protocol (STP) reconverges because the STP bridge protocol data unit (BPDU) cannot be sent or received.
- CLI is slow or unresponsive.

Using protocol storm protection, you can control the rate at which control packets are sent to the switch by specifying the upper threshold for the packet flow rate. The supported protocols are ARP, ARP snooping, Dynamic Host Configuration Protocol (DHCP) v4, DHCP snooping, Internet Group Management Protocol (IGMP), and IGMP snooping.

When the packet rate exceeds the defined threshold, the switch drops all traffic arriving on the specified virtual port for 30 seconds. The packet rate is measured again, and protocol storm protection is again applied if necessary.

For further protection, you can manually error disable the virtual port, blocking all incoming traffic on the virtual port. You can manually enable the virtual port or set a time interval for automatic re-enabling of the virtual port.



Note

Excess packets are dropped on no more than two virtual ports.
Virtual port error disabling is not supported for EtherChannel and Flexlink interfaces.

Default Protocol Storm Protection Configuration

Protocol storm protection is disabled by default. When it is enabled, auto-recovery of the virtual port is disabled by default.

Enabling Protocol Storm Protection

Beginning in privileged EXEC mode, follow these steps to configure protocol storm protection.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	psp {arp dhcp igmp} pps value	Configure protocol storm protection for ARP, IGMP, or DHCP. For <i>value</i> , specify the threshold value for the number of packets per second. If the traffic exceeds this value, protocol storm protection is enforced. The range is from 5 to 50 packets per second.
Step 3	errdisable detect cause psp	(Optional) Enable error-disable detection for protocol storm protection. If this feature is enabled, the virtual port is error disabled. If this feature is disabled, the port drops excess packets without error disabling the port.
Step 4	errdisable recovery interval time	(Optional) Configure an auto-recovery time (in seconds) for error-disabled virtual ports. When a virtual port is error-disabled, the switch auto-recovers after this time. The range is from 30 to 86400 seconds.
Step 5	end	Return to privileged EXEC mode.
Step 6	show psp config {arp dhcp igmp}	Verify your entries.

This example shows how to configure protocol storm protection to drop incoming DHCP traffic on DHCP when it exceeds 35 packets per second.

```
Switch# configure terminal
Switch(config)# psp dhcp pps 35
```

To disable protocol storm protection for a specific protocol, use the **no psp {arp | dhcp | igmp}** privileged EXEC command.

To disable error-disable detection for protocol storm protection, use the **no errdisable detect cause psp** global configuration command.

To manually re-enable an error-disabled virtual port, use the **errdisable recovery cause psp** global configuration command.

To disable auto-recovery of error-disabled ports, use the **no errdisable recovery cause psp** global configuration command.

When protocol storm protection is configured, a counter records the number of dropped packets. To see this counter, use the **show psp statistics [arp | igmp | dhcp]** privileged EXEC command. To clear the counter for a protocol, use the **clear psp counter [arp | igmp | dhcp]** 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 storm-control** and **show port-security** privileged EXEC commands display those storm control and port security settings.

To display traffic control information, use one or more of the privileged EXEC commands in [Table 1-4](#).

Table 1-4 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 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.



CHAPTER 1

Configuring UDLD

This chapter describes how to configure the UniDirectional Link Detection (UDLD) protocol on the Catalyst 2960 and 2960-S switches. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.



Note

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release.

This chapter consists of these sections:

- [Understanding UDLD, page 1-1](#)
- [Configuring UDLD, page 1-3](#)
- [Displaying UDLD Status, page 1-7](#)

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 disables the affected port 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 can also detect unidirectional links due to one-way traffic on fiber-optic and twisted-pair links and to misconnected ports on fiber-optic links.

In normal and aggressive modes, 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 ports. 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 detects 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 port.

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 detects 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 detects a unidirectional link by using the previous detection methods. UDLD in aggressive mode can also detect a unidirectional link on a point-to-point link on which no failure between the two devices is allowed. It can also detect a unidirectional link when one of these problems exists:

- On fiber-optic or twisted-pair links, one of the ports cannot send or receive traffic.
- On fiber-optic or twisted-pair links, one of the ports is down while the other is up.
- One of the fiber strands in the cable is disconnected.

In these cases, UDLD disables the affected port.

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 fiber strands in a cable are working normally from a Layer 1 perspective, UDLD in aggressive mode detects whether those fiber strands are connected correctly and whether traffic is flowing bidirectionally between the correct neighbors. This check cannot be performed by autonegotiation because autonegotiation operates at Layer 1.

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 port 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 a port is disabled and UDLD is running, whenever UDLD is disabled on a port, or whenever the switch is reset, UDLD clears all existing cache entries for the ports 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 be shut down. When UDLD is in aggressive mode, the link is considered unidirectional, and the port is disabled.

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 a port 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 port if, after the fast train of messages, the link state is still undetermined.

Figure 1-1 shows an example of a unidirectional link condition.

Figure 1-1 UDLD Detection of a Unidirectional Link



84986

Configuring UDLD

These sections contain this configuration information:

- [Default UDLD Configuration, page 1-4](#)
- [Configuration Guidelines, page 1-4](#)
- [Enabling UDLD Globally, page 1-5](#)
- [Enabling UDLD on an Interface, page 1-6](#)
- [Resetting an Interface Disabled by UDLD, page 1-6](#)

Default UDLD Configuration

Table 1-1 shows the default UDLD configuration.

Table 1-1 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

Configuration Guidelines

These are the UDLD configuration guidelines:

- UDLD is not supported on ATM ports.
- A UDLD-capable port cannot detect a unidirectional link if it is connected to a UDLD-incapable port of another switch.
- When configuring the mode (normal or aggressive), make sure that the same mode is configured on both sides of the link.



Caution

Loop guard works only on point-to-point links. We recommend that each end of the link has a directly connected device that is running STP.

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 and all members in the switch stack:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	udld { aggressive enable message time message-timer-interval }	<p>Specify the UDLD mode of operation:</p> <ul style="list-style-type: none"> • 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 1-1. • message time message-timer-interval—Configures the period of time between UDLD probe messages on ports that are in the advertisement phase and are detected to be bidirectional. The range is from 7 to 90 seconds. The default value is 15. <p>Note The global UDLD setting is automatically applied to switches that join the switch stack.</p> <p>Note This command affects fiber-optic ports only. Use the udld interface configuration command to enable UDLD on other port types. For more information, see the “Enabling UDLD on an Interface” section on page 1-6.</p>
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 a port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port to be enabled for UDLD, and enter interface configuration mode.
Step 3	udld port [aggressive]	UDLD is disabled by default. Note When a switch joins a switch stack, it retains its interface-specific UDLD settings. <ul style="list-style-type: none"> • udld port—Enables UDLD in normal mode on the specified port. • udld port aggressive—Enables UDLD in aggressive mode on the specified port. Note Use the no udld port interface configuration command to disable UDLD on a specified fiber-optic port. For more information about aggressive and normal modes, see the “Modes of Operation” section on page 1-1 .
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 ports disabled by UDLD:

	Command	Purpose
Step 1	udld reset	Reset all ports disabled by UDLD.
Step 2	show udld	Verify your entries.

You can also bring up the port by using these commands:

- The **shutdown** interface configuration command followed by the **no shutdown** interface configuration command restarts the disabled port.
- The **no udld {aggressive | enable}** global configuration command followed by the **udld {aggressive | enable}** global configuration command re-enables the disabled ports.
- The **no udld port** 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 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 port or for all ports, use the **show udld** [*interface-id*] privileged EXEC command.

For detailed information about the fields in the command output, see the command reference for this release.



CHAPTER 1

Configuring CDP

This chapter describes how to configure Cisco Discovery Protocol (CDP) on the Catalyst 2960 and 2960-S switches. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.



Note

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release and the “System Management Commands” section in the *Cisco IOS Configuration Fundamentals Command Reference, Release 12.4*.

This chapter consists of these sections:

- [Understanding CDP, page 1-1](#)
- [Configuring CDP, page 1-2](#)
- [Monitoring and Maintaining CDP, page 1-5](#)

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.

On the switch, CDP enables Network Assistant to display a graphical view of the network. The switch uses CDP to find cluster candidates and maintain information about cluster members and other devices up to three cluster-enabled devices away from the command switch by default.

For a switch and connected endpoint devices running Cisco Medianet

- CDP identifies connected endpoints that communicate directly with the switch.
- To prevent duplicate reports of neighboring devices, only one wired switch reports the location information.
- The wired switch and the endpoints both send and receive location information.

For information, go to:

http://www.cisco.com/en/US/docs/ios/netmgmt/configuration/guide/nm_cdp_discover.html.

The switch supports CDP Version 2.

CDP and Switch Stacks

A switch stack appears as a single switch in the network. Therefore, CDP discovers the switch stack, not the individual stack members. The switch stack sends CDP messages to neighboring network devices when there are changes to the switch stack membership, such as stack members being added or removed.

Configuring CDP

These sections contain this configuration information:

- [Default CDP Configuration, page 1-2](#)
- [Configuring the CDP Characteristics, page 1-3](#)
- [Disabling and Enabling CDP, page 1-3](#)
- [Disabling and Enabling CDP on an Interface, page 1-4](#)

Default CDP Configuration

Table 1-1 shows the default CDP configuration.

Table 1-1 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.
Step 6	show cdp	Verify your settings.
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 CDP characteristics.

```
Switch# configure terminal
Switch(config)# cdp timer 50
Switch(config)# cdp holdtime 120
Switch(config)# cdp advertise-v2
Switch(config)# end
```

For additional CDP **show** commands, see the [“Monitoring and Maintaining CDP”](#) section on page 1-5.

Disabling and Enabling CDP

CDP is enabled by default.



Note Switch clusters and other Cisco devices (such as Cisco IP Phones) regularly exchange CDP messages. Disabling CDP can interrupt cluster discovery and device connectivity. For more information, see [Chapter 1, “Clustering Switches”](#) and see *Getting Started with Cisco Network Assistant*, available on Cisco.com.

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 to receive CDP information.

Beginning in privileged EXEC mode, follow these steps to disable CDP on a port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface on which you are disabling CDP, and enter interface configuration mode.
Step 3	no cdp enable	Disable CDP on the 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 a port when it has been disabled:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface on which you are enabling CDP, and enter interface configuration mode.
Step 3	cdp enable	Enable CDP on the 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 gigabitethernet0/1
Switch(config-if)# cdp enable
Switch(config-if)# end
```



Note

Voice VLAN is not counted against port security when CDP is disabled on the switch interface.

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. You can limit the display to the interface about which you want information.

Command	Description
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 of 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.



CHAPTER 1

Configuring LLDP, LLDP-MED, and Wired Location Service

This chapter describes how to configure the Link Layer Discovery Protocol (LLDP), LLDP Media Endpoint Discovery (LLDP-MED) and wired location service on the Catalyst 2960 and 2960-S switches.



Note

To use wired location service, the switch must be running the LAN Base image.

Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.



Note

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release and the “System Management Commands” section in the *Cisco IOS Configuration Fundamentals Command Reference, Release 12.4*.

- [Understanding LLDP, LLDP-MED, and Wired Location Service, page 1-1](#)
- [Configuring LLDP, LLDP-MED, and Wired Location Service, page 1-5](#)
- [Monitoring and Maintaining LLDP, LLDP-MED, and Wired Location Service, page 1-11](#)

Understanding LLDP, LLDP-MED, and Wired Location Service

LLDP

The Cisco Discovery Protocol (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). CDP allows network management applications to automatically discover and learn about other Cisco devices connected to the network.

To support non-Cisco devices and to allow for interoperability between other devices, the switch supports the IEEE 802.1AB Link Layer Discovery Protocol (LLDP). LLDP is a neighbor discovery protocol that is used for network devices to advertise information about themselves to other devices on the network. This protocol runs over the data-link layer, which allows two systems running different network layer protocols to learn about each other.

LLDP supports a set of attributes that it uses to discover neighbor devices. These attributes contain type, length, and value descriptions and are referred to as TLVs. LLDP supported devices can use TLVs to receive and send information to their neighbors. This protocol can advertise details such as configuration information, device capabilities, and device identity.

The switch supports these basic management TLVs. These are mandatory LLDP TLVs.

- Port description TLV
- System name TLV
- System description TLV
- System capabilities TLV
- Management address TLV

These organizationally specific LLDP TLVs are also advertised to support LLDP-MED.

- Port VLAN ID TLV ((IEEE 802.1 organizationally specific TLVs)
- MAC/PHY configuration/status TLV(IEEE 802.3 organizationally specific TLVs)



Note

A switch stack appears as a single switch in the network. Therefore, LLDP discovers the switch stack, not the individual stack members.

When you configure LLDP or CDP location information on a per-port basis, remote devices can send Cisco Medianet location information to the switch. For information, go to http://www.cisco.com/en/US/docs/ios/netmgmt/configuration/guide/nm_cdp_discover.html.

LLDP-MED

LLDP for Media Endpoint Devices (LLDP-MED) is an extension to LLDP that operates between endpoint devices such as IP phones and network devices such as switches. It specifically provides support for voice over IP (VoIP) applications and provides additional TLVs for capabilities discovery, network policy, Power over Ethernet, inventory management and location information. By default, all LLDP-MED TLVs are enabled.

LLDP-MED supports these TLVs:

- LLDP-MED capabilities TLV

Allows LLDP-MED endpoints to determine the capabilities that the connected device supports and has enabled.
- Network policy TLV

Allows both network connectivity devices and endpoints to advertise VLAN configurations and associated Layer 2 and Layer 3 attributes for the specific application on that port. For example, the switch can notify a phone of the VLAN number that it should use. The phone can connect to any switch, obtain its VLAN number, and then start communicating with the call control.

By defining a network-policy profile TLV, you can create a profile for voice and voice-signalling by specifying the values for VLAN, class of service (CoS), differentiated services code point (DSCP), and tagging mode. These profile attributes are then maintained centrally on the switch and propagated to the phone.

- Power management TLV

Enables advanced power management between LLDP-MED endpoint and network connectivity devices. Allows switches and phones to convey power information, such as how the device is powered, power priority, and how much power the device needs.

LLDP-MED also supports an extended power TLV to advertise fine-grained power requirements, end-point power priority, and end-point and network connectivity-device power status.

Starting with Cisco IOS Release 12.2(52)SE, when LLDP is enabled and power is applied to a port, the power TLV determines the actual power requirement of the endpoint device so that the system power budget can be adjusted accordingly. The switch processes the requests and either grants or denies power based on the current power budget. If the request is granted, the switch updates the power budget. If the request is denied, the switch turns off power to the port, generates a syslog message, and updates the power budget. If LLDP-MED is disabled or if the endpoint does not support the LLDP-MED power TLV, the initial allocation value (15.4 W) is used throughout the duration of the connection.

You can change power settings by entering the **power inline {auto [max *max-wattage*] | never | static [max *max-wattage*]}** interface configuration command. By default the PoE interface is in **auto** mode; If no value is specified, the maximum is allowed (15.4 W).

- Inventory management TLV

Allows an endpoint to send detailed inventory information about itself to the switch, including information hardware revision, firmware version, software version, serial number, manufacturer name, model name, and asset ID TLV.

- Location TLV

Provides location information from the switch to the endpoint device. The location TLV can send this information:

- Civic location information

Provides the civic address information and postal information. Examples of civic location information are street address, road name, and postal community name information.

- ELIN location information

Provides the location information of a caller. The location is determined by the Emergency location identifier number (ELIN), which is a phone number that routes an emergency call to the local public safety answering point (PSAP) and which the PSAP can use to call back the emergency caller.

Wired Location Service



Note

To use wired location service, the switch must be running the LAN Base image.

The switch uses the wired location service feature to send location and attachment tracking information for its connected devices to a Cisco Mobility Services Engine (MSE). The tracked device can be a wireless endpoint, a wired endpoint, or a wired switch or controller. The switch notifies the MSE of device link up and link down events through the Network Mobility Services Protocol (NMSP) location and attachment notifications.

The MSE starts the NMSP connection to the switch, which opens a server port. When the MSE connects to the switch there are a set of message exchanges to establish version compatibility and service exchange information followed by location information synchronization. After connection, the switch periodically sends location and attachment notifications to the MSE. Any link up or link down events detected during an interval are aggregated and sent at the end of the interval.

When the switch determines the presence or absence of a device on a link-up or link-down event, it obtains the client-specific information such as the MAC address, IP address, and username. If the client is LLDP-MED- or CDP-capable, the switch obtains the serial number and UDI through the LLDP-MED location TLV or CDP.

Depending on the device capabilities, the switch obtains this client information at link up:

- Slot and port specified in port connection
- MAC address specified in the client MAC address
- IP address specified in port connection
- 802.1X username if applicable
- Device category is specified as a *wired station*
- State is specified as *new*
- Serial number, UDI
- Model number
- Time in seconds since the switch detected the association

Depending on the device capabilities, the switch obtains this client information at link down:

- Slot and port that was disconnected
- MAC address
- IP address
- 802.1X username if applicable
- Device category is specified as a *wired station*
- State is specified as *delete*
- Serial number, UDI
- Time in seconds since the switch detected the disassociation

When the switch shuts down, it sends an attachment notification with the state *delete* and the IP address before closing the NMSP connection to the MSE. The MSE interprets this notification as disassociation for all the wired clients associated with the switch.

If you change a location address on the switch, the switch sends an NMSP location notification message that identifies the affected ports and the changed address information.

Configuring LLDP, LLDP-MED, and Wired Location Service

- [Default LLDP Configuration, page 1-5](#)
- [Configuration Guidelines, page 1-5](#)
- [Enabling LLDP, page 1-6](#)
- [Configuring LLDP Characteristics, page 1-6](#)
- [Configuring LLDP-MED TLVs, page 1-7](#)
- [Configuring Network-Policy TLV, page 1-8](#)
- [Configuring Location TLV and Wired Location Service, page 1-10](#)

Default LLDP Configuration

Table 1-1 Default LLDP Configuration

Feature	Default Setting
LLDP global state	Disabled
LLDP holdtime (before discarding)	120 seconds
LLDP timer (packet update frequency)	30 seconds
LLDP reinitialization delay	2 seconds
LLDP tlv-select	Disabled to send and receive all TLVs
LLDP interface state	Disabled
LLDP receive	Disabled
LLDP transmit	Disabled
LLDP med-tlv-select	Disabled to send all LLDP-MED TLVs. When LLDP is globally enabled, LLDP-MED-TLV is also enabled.

Configuration Guidelines

- If the interface is configured as a tunnel port, LLDP is automatically disabled.
- If you first configure a network-policy profile on an interface, you cannot apply the **switchport voice vlan** command on the interface. If the **switchport voice vlan *vlan-id*** is already configured on an interface, you can apply a network-policy profile on the interface. This way the interface has the voice or voice-signaling VLAN network-policy profile applied on the interface.
- You cannot configure static secure MAC addresses on an interface that has a network-policy profile.
- You cannot configure a network-policy profile on a private-VLAN port.
- For wired location to function, you must first enter the **ip device tracking** global configuration command.

Enabling LLDP

Beginning in privileged EXEC mode, follow these steps to enable LLDP:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	lldp run	Enable LLDP globally on the switch.
Step 3	interface <i>interface-id</i>	Specify the interface on which you are enabling LLDP, and enter interface configuration mode.
Step 4	lldp transmit	Enable the interface to send LLDP packets.
Step 5	lldp receive	Enable the interface to receive LLDP packets.
Step 6	end	Return to privileged EXEC mode.
Step 7	show lldp	Verify the configuration.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable LLDP, use the **no lldp run** global configuration command. To disable LLDP on an interface, use the **no lldp transmit** and the **no lldp receive** interface configuration commands.

This example shows how to globally enable LLDP.

```
Switch# configure terminal
Switch(config)# lldp run
Switch(config)# end
```

This example shows how to enable LLDP on an interface.

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# lldp transmit
Switch(config-if)# lldp receive
Switch(config-if)# end
```

Configuring LLDP Characteristics

You can configure the frequency of LLDP updates, the amount of time to hold the information before discarding it, and the initialization delay time. You can also select the LLDP and LLDP-MED TLVs to send and receive.

Beginning in privileged EXEC mode, follow these steps to configure the LLDP characteristics.



Note Steps 2 through 5 are optional and can be performed in any order.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	lldp holdtime <i>seconds</i>	(Optional) Specify the amount of time a receiving device should hold the information from your device before discarding it. The range is 0 to 65535 seconds; the default is 120 seconds.
Step 3	lldp reinit <i>delay</i>	(Optional) Specify the delay time in seconds for LLDP to initialize on an interface. The range is 2 to 5 seconds; the default is 2 seconds.
Step 4	lldp timer <i>rate</i>	(Optional) Set the sending frequency of LLDP updates in seconds. The range is 5 to 65534 seconds; the default is 30 seconds.
Step 5	lldp tlv-select	(Optional) Specify the LLDP TLVs to send or receive.
Step 6	interface <i>interface-id</i>	Specify the interface on which you are enabling LLDP, and enter interface configuration mode.
Step 7	lldp med-tlv-select	(Optional) Specify the LLDP-MED TLVs to send or receive.
Step 8	end	Return to privileged EXEC mode.
Step 9	show lldp	Verify the configuration.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no** form of each of the LLDP commands to return to the default setting.

This example shows how to configure LLDP characteristics.

```
Switch# configure terminal
Switch(config)# lldp holdtime 120
Switch(config)# lldp reinit 2
Switch(config)# lldp timer 30
Switch(config)# end
```

Configuring LLDP-MED TLVs

By default, the switch only sends LLDP packets until it receives LLDP-MED packets from the end device. It then sends LLDP packets with MED TLVs, as well. When the LLDP-MED entry has been aged out, it again only sends LLDP packets.

By using the **lldp** interface configuration command, you can configure the interface not to send the TLVs listed in [Table 1-2](#).

Table 1-2 LLDP-MED TLVs

LLDP-MED TLV	Description
inventory-management	LLDP-MED inventory management TLV
location	LLDP-MED location TLV

Table 1-2 LLDP-MED TLVs

LLDP-MED TLV	Description
network-policy	LLDP-MED network policy TLV
power-management	LLDP-MED power management TLV

Beginning in privileged EXEC mode, follow these steps to enable a TLV on an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface on which you are configuring an LLDP-MED TLV, and enter interface configuration mode.
Step 3	lldp med-tlv-select <i>tlv</i>	Specify the TLV to enable.
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 a TLV on an interface:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# lldp med-tlv-select inventory-management
Switch(config-if)# end
```

Configuring Network-Policy TLV

Beginning in privileged EXEC mode, follow these steps to create a network-policy profile, configure the policy attributes, and apply it to an interface.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	network-policy profile <i>profile number</i>	Specify the network-policy profile number, and enter network-policy configuration mode. The range is 1 to 4294967295.

	Command	Purpose
Step 3	{ voice voice-signaling } vlan [vlan-id { cos cvalue dscp dvalue }] [[dot1p { cos cvalue dscp dvalue }] none untagged]	Configure the policy attributes: voice —Specify the voice application type. voice-signaling —Specify the voice-signaling application type. vlan —Specify the native VLAN for voice traffic. <i>vlan-id</i> —(Optional) Specify the VLAN for voice traffic. The range is 1 to 4094. cos cvalue —(Optional) Specify the Layer 2 priority class of service (CoS) for the configured VLAN. The range is 0 to 7; the default is 5. dscp dvalue —(Optional) Specify the differentiated services code point (DSCP) value for the configured VLAN. The range is 0 to 63; the default is 46. dot1p —(Optional) Configure the telephone to use IEEE 802.1p priority tagging and use VLAN 0 (the native VLAN). none —(Optional) Do not instruct the IP telephone about the voice VLAN. The telephone uses the configuration from the telephone key pad. untagged —(Optional) Configure the telephone to send untagged voice traffic. This is the default for the telephone.
Step 4	exit	Return to global configuration mode.
Step 5	interface interface-id	Specify the interface on which you are configuring a network-policy profile, and enter interface configuration mode.
Step 6	network-policy profile number	Specify the network-policy profile number.
Step 7	lldp med-tlv-select network-policy	Specify the network-policy TLV.
Step 8	end	Return to privileged EXEC mode.
Step 9	show network-policy profile	Verify the configuration.
Step 10	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 setting.

This example shows how to configure VLAN 100 for voice application with CoS and to enable the network-policy profile and network-policy TLV on an interface:

```
Switch# configure terminal
Switch(config)# network-policy 1
Switch(config-network-policy)# voice vlan 100 cos 4
Switch(config-network-policy)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# network-policy profile 1
Switch(config-if)# lldp med-tlv-select network-policy
```

This example shows how to configure the voice application type for the native VLAN with priority tagging:

```
Switch(config-network-policy)# voice vlan dot1p cos 4
Switch(config-network-policy)# voice vlan dot1p dscp 34
```

Configuring Location TLV and Wired Location Service



Note To use wired location service, the switch must be running the LAN Base image.

Beginning in privileged EXEC mode, follow these steps to configure location information for an endpoint and to apply it to an interface.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	location { admin-tag <i>string</i> civic-location identifier <i>id</i> elin-location <i>string identifier id</i> }	Specify the location information for an endpoint. <ul style="list-style-type: none"> • admin-tag—Specify an administrative tag or site information. • civic-location—Specify civic location information. • elin-location—Specify emergency location information (ELIN). • identifier id—Specify the ID for the civic location. • <i>string</i>—Specify the site or location information in alphanumeric format.
Step 3	exit	Return to global configuration mode.
Step 4	interface <i>interface-id</i>	Specify the interface on which you are configuring the location information, and enter interface configuration mode.
Step 5	location { additional-location-information <i>word</i> civic-location-id <i>id</i> elin-location-id <i>id</i> }	Enter location information for an interface: <ul style="list-style-type: none"> additional-location-information—Specify additional information for a location or place. civic-location-id—Specify global civic location information for an interface. elin-location-id—Specify emergency location information for an interface. <i>id</i>—Specify the ID for the civic location or the ELIN location. The ID range is 1 to 4095. <i>word</i>—Specify a word or phrase with additional location information.
Step 6	end	Return to privileged EXEC mode.
Step 7	show location admin-tag <i>string</i> or show location civic-location identifier <i>id</i> or show location elin-location identifier <i>id</i>	Verify the configuration.
Step 8	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 setting.

This example shows how to configure civic location information on the switch:

```
Switch(config)# location civic-location identifier 1
Switch(config-civic)# number 3550
Switch(config-civic)# primary-road-name "Cisco Way"
Switch(config-civic)# city "San Jose"
Switch(config-civic)# state CA
Switch(config-civic)# building 19
Switch(config-civic)# room C6
Switch(config-civic)# county "Santa Clara"
Switch(config-civic)# country US
Switch(config-civic)# end
```

Beginning in privileged EXEC mode, follow these steps to enable wired location service on the switch.



Note

Your switch must be running the cryptographic (encrypted) software image to enable the **nmsp** global configuration commands.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	nmsp enable	Enable the NMSP features on the switch.
Step 3	nmsp notification interval {attachment location} interval-seconds	Specify the NMSP notification interval. attachment —Specify the attachment notification interval. location —Specify the location notification interval. <i>interval-seconds</i> —Duration in seconds before the switch sends the MSE the location or attachment updates. The range is 1 to 30; the default is 30.
Step 4	end	Return to privileged EXEC mode.
Step 5	show network-policy profile	Verify the configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to enable NMSP on a switch and to set the location notification time to 10 seconds:

```
Switch(config)# nmsp enable
Switch(config)# nmsp notification interval location 10
```

Monitoring and Maintaining LLDP, LLDP-MED, and Wired Location Service

To monitor and maintain LLDP, LLDP-MED, and wired location service on your device, perform one or more of these tasks, beginning in privileged EXEC mode.

Command	Description
clear lldp counters	Reset the traffic counters to zero.
clear lldp table	Delete the LLDP neighbor information table.

Command	Description
clear nmsp statistics	Clear the NMSP statistic counters.
show lldp	Display global information, such as frequency of transmissions, the holdtime for packets being sent, and the delay time before LLDP initializes on an interface.
show lldp entry <i>entry-name</i>	Display information about a specific neighbor. You can enter an asterisk (*) to display all neighbors, or you can enter the neighbor name.
show lldp interface [<i>interface-id</i>]	Display information about interfaces with LLDP enabled. You can limit the display to a specific interface.
show lldp neighbors [<i>interface-id</i>] [detail]	Display information about neighbors, including device type, interface type and number, holdtime settings, capabilities, and port ID. You can limit the display to neighbors of a specific interface or expand the display for more detailed information.
show lldp traffic	Display LLDP counters, including the number of packets sent and received, number of packets discarded, and number of unrecognized TLVs.
show location admin-tag <i>string</i>	Display the location information for the specified administrative tag or site.
show location civic-location identifier <i>id</i>	Display the location information for a specific global civic location.
show location elin-location identifier <i>id</i>	Display the location information for an emergency location.
show network-policy profile	Display the configured network-policy profiles.
show nmsp	Display the NMSP information.



CHAPTER 1

Configuring SPAN and RSPAN



Note

To use RSPAN, the switch must be running the LAN Base image.

This chapter describes how to configure Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) on the Catalyst 2960 and 2960-S switches. Unless otherwise noted, the term *switch* refers to a standalone switch and a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.



Note

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release.

- [Understanding SPAN and RSPAN, page 1-1](#)
- [Configuring SPAN and RSPAN, page 1-10](#)
- [Displaying SPAN and RSPAN Status, page 1-23](#)

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.


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.

- [Local SPAN, page 1-2](#)
- [Remote SPAN, page 1-3](#)
- [SPAN and RSPAN Concepts and Terminology, page 1-4](#)
- [SPAN and RSPAN Interaction with Other Features, page 1-9](#)

Local SPAN

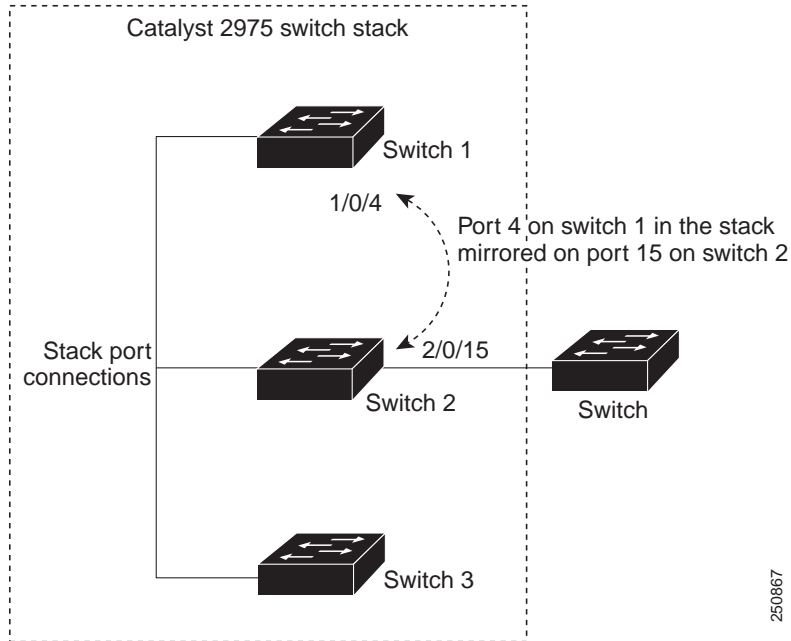
Local SPAN supports a SPAN session entirely within one switch; all source ports or source VLANs and destination ports are in the same switch or switch stack. Local SPAN copies traffic from one or more source ports in any VLAN or from one or more VLANs to a destination port for analysis. For example, in [Figure 1-1](#), 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 1-1 Example of Local SPAN Configuration on a Single Switch



[Figure 1-2](#) is an example of a local SPAN in a switch stack, where the source and destination ports reside on different stack members.

Figure 1-2 Example of Local SPAN Configuration on a Switch Stack



Remote SPAN

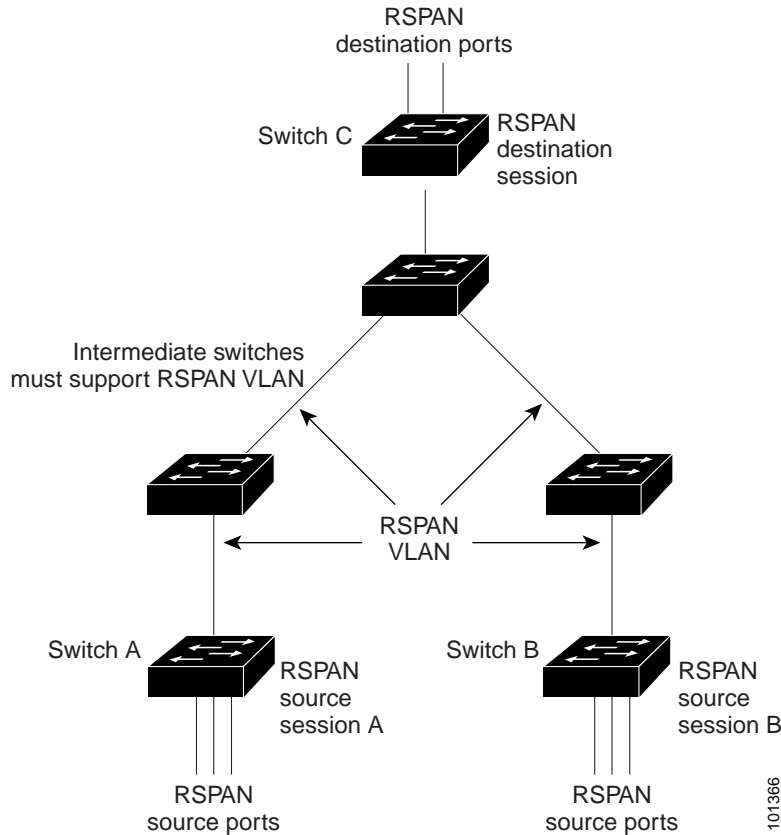


Note

To use RSPAN, the switch must be running the LAN Base image.

RSPAN supports source ports, source VLANs, and destination ports on different switches (or different switch stacks), enabling remote monitoring of multiple switches across your network. Figure 1-3 shows source ports on Switch A and Switch B. 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, as shown on Switch C in the figure.

Figure 1-3 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 respond to the requirements of the RSPAN VLAN (see the “[RSPAN VLAN](#)” section on page 1-9).

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.
- With the LAN Base license, the switch supports up to two source sessions (local SPAN and RSPAN source sessions). You can run both a local SPAN and an RSPAN source session in the same switch stack. The switch stack supports a total of 64 source and RSPAN destination sessions.
- With the LAN Lite license, the switch supports configuration of only one SPAN session.
- You can have multiple destination ports in a SPAN session, but no more than 64 destination ports per switch stack.
- You can configure two separate SPAN or RSPAN source sessions with separate or overlapping sets of SPAN source ports and VLANs.
- SPAN sessions do not interfere with the normal operation of the switch. However, an oversubscribed SPAN destination, for example, a 10-Mb/s port monitoring a 100-Mb/s 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 stack.

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, 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.

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 or IEEE 802.1Q—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 and IEEE 802.1Q 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.

Source Ports

A source port (also called a *monitored port*) is a switched 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, or voice VLAN 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 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 stack 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 or switch stack 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.



Note

Exception. When QoS is configured on the SPAN destination port, QoS takes effect immediately.

- If the port was in an EtherChannel group, it is removed from the group while it is a destination 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 stack 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 or IEEE 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 or IEEE 802.1Q-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:

- 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.
- 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.
- 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.

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 *suspended* 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 IEEE 802.1x port can be a SPAN source port. You can enable IEEE 802.1x on a port that is a SPAN destination port; however, IEEE 802.1x is disabled until the port is removed as a SPAN destination.

For SPAN sessions, do not enable IEEE 802.1x on ports with monitored egress when ingress forwarding is enabled on the destination port. For RSPAN source sessions, do not enable IEEE 802.1x on any ports that are egress monitored.

SPAN and RSPAN and Switch Stacks

Because the stack of switches is treated as one logical switch, local SPAN source ports and destination ports can be in different switches in the stack. Therefore, the addition or deletion of switches in the stack can affect a local SPAN session, as well as an RSPAN source or destination session. An active session can become inactive when a switch is removed from the stack or an inactive session can become active when a switch is added to the stack.

For more information about switch stacks, see [Chapter 1, “Managing Switch Stacks.”](#)

Configuring SPAN and RSPAN

- [Default SPAN and RSPAN Configuration, page 1-10](#)
- [Configuring Local SPAN, page 1-11](#)
- [Configuring RSPAN, page 1-17](#)

Default SPAN and RSPAN Configuration

[Table 1-1](#) shows the default SPAN and RSPAN configuration.

Table 1-1 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

Table 1-1 Default SPAN and RSPAN Configuration (continued)

Feature	Default Setting
VLAN filtering	On a trunk interface used as a source port, all VLANs are monitored.
RSPAN VLANs	None configured.

Configuring Local SPAN

- [SPAN Configuration Guidelines, page 1-11](#)
- [Creating a Local SPAN Session, page 1-12](#)
- [Creating a Local SPAN Session and Configuring Incoming Traffic, page 1-14](#)
- [Specifying VLANs to Filter, page 1-16](#)

SPAN Configuration Guidelines

- 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 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.
- With the LAN Lite license, the switch supports configuration of only one 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 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 1 to 66. For <i>interface-id</i> , specify the source port or source VLAN to monitor. <ul style="list-style-type: none"> 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 6. For <i>vlan-id</i>, specify the source VLAN to monitor. The range is 1 to 4094 (excluding the RSPAN VLAN). <p>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.</p> <p>(Optional) [, -] Specify a series or range of interfaces. Enter a space before and after the comma; enter a space before and after the hyphen.</p> <p>(Optional) Specify the direction of traffic to monitor. If you do not specify a traffic direction, the SPAN monitors both sent and received traffic.</p> <ul style="list-style-type: none"> both—Monitor both received and sent traffic. This is the default. rx—Monitor received traffic. tx—Monitor sent traffic. <p>Note You can use the monitor session <i>session_number</i> source command multiple times to configure multiple source ports.</p>

	Command	Purpose
Step 4	monitor session <i>session_number</i> destination { interface <i>interface-id</i> [, -] [encapsulation { dot1q 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 dot1q to specify that the destination interface uses the IEEE 802.1Q encapsulation method. (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 options are ignored with the **no** form of the command.

This example shows how to set up 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 gigabitethernet0/1
Switch(config)# monitor session 1 destination interface gigabitethernet0/2
encapsulation replicate
Switch(config)# end
```

This example shows how to remove port 1 as a SPAN source for SPAN session 1:

```
Switch(config)# no monitor session 1 source interface gigabitethernet0/1
Switch(config)# end
```

This example shows how to disable received traffic monitoring on port 1, which was configured for bidirectional monitoring:

```
Switch(config)# no monitor session 1 source interface gigabitethernet0/1 rx
```

The monitoring of traffic received on port 1 is disabled, 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 destination Gigabit Ethernet port 2. 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 gigabitethernet0/2
Switch(config)# monitor session 2 source vlan 10
Switch(config)# end
```

Creating a Local SPAN Session and Configuring Incoming 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 incoming traffic on the destination port for a network security device (such as a Cisco IDS Sensor Appliance).

For details about the keywords not related to incoming traffic, see the [“Creating a Local SPAN Session” section on page 1-12](#).

	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 { dot1q replicate }] [ingress { dot1q vlan <i>vlan-id</i> 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 dot1q to specify that the destination interface uses the IEEE 802.1Q encapsulation method. (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 forwarding of incoming traffic on the destination port and to specify the encapsulation type: <ul style="list-style-type: none"> • dot1q vlan <i>vlan-id</i>—Accept incoming packets with IEEE 802.1Q encapsulation with the specified VLAN as the default VLAN. • untagged vlan <i>vlan-id</i> or vlan <i>vlan-id</i>—Accept incoming 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 to enable ingress forwarding with IEEE 802.1Q encapsulation and VLAN 6 as the default ingress VLAN.

```
Switch(config)# no monitor session 2
Switch(config)# monitor session 2 source gigabitethernet0/1 rx
Switch(config)# monitor session 2 destination interface gigabitethernet0/2 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 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 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 { dot1q 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 dot1q to specify that the destination interface uses the IEEE 802.1Q encapsulation method. (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 VLAN 9 to destination Gigabit Ethernet port 1.

```
Switch(config)# no monitor session 2
Switch(config)# monitor session 2 source interface gigabitethernet0/2 rx
Switch(config)# monitor session 2 filter vlan 1 - 5, 9
Switch(config)# monitor session 2 destination interface gigabitethernet0/1
Switch(config)# end
```

Configuring RSPAN

- [RSPAN Configuration Guidelines, page 1-17](#)
- [Configuring a VLAN as an RSPAN VLAN, page 1-18](#)
- [Creating an RSPAN Source Session, page 1-18](#)
- [Creating an RSPAN Destination Session, page 1-20](#)
- [Creating an RSPAN Destination Session and Configuring Incoming Traffic, page 1-21](#)
- [Specifying VLANs to Filter, page 1-22](#)

RSPAN Configuration Guidelines

- All the items in the “[SPAN Configuration Guidelines](#)” section on page 1-11 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 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 2 to 1001 and 1006 to 4094. 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 1 to 66. Specify all to remove all RSPAN sessions, local to remove all local sessions, or remote to remove all remote SPAN sessions.

Command	Purpose
Step 3 monitor session <i>session_number</i> source { interface <i>interface-id</i> vlan <i>vlan-id</i> } [, -] [both rx tx]	<p>Specify the RSPAN session and the source port (monitored port). For <i>session_number</i>, the range is 1 to 66.</p> <p>Enter a source port or source VLAN for the RSPAN session:</p> <ul style="list-style-type: none"> 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 6. For <i>vlan-id</i>, specify the source VLAN to monitor. The range is 1 to 4094 (excluding the RSPAN VLAN). <p>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.</p> <p>(Optional) [, -] Specify a series or range of interfaces. Enter a space before and after the comma; enter a space before and after the hyphen.</p> <p>(Optional) Specify the direction of traffic to monitor. If you do not specify a traffic direction, the source interface sends both sent and received traffic.</p> <ul style="list-style-type: none"> both—Monitor both received and sent traffic. rx—Monitor received traffic. tx—Monitor sent traffic.
Step 4 monitor session <i>session_number</i> destination remote vlan <i>vlan-id</i>	<p>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.</p>
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 gigabitethernet0/1 tx
Switch(config)# monitor session 1 source interface gigabitethernet0/2 rx

Switch(config)# monitor session 1 source interface port-channel 2
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 or switch stack; that is, not the switch or switch stack 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. 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 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 1 to 66. For <i>vlan-id</i> , specify the source RSPAN VLAN to monitor.
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. 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. 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 port 1 as the destination interface:

```
Switch(config)# monitor session 1 source remote vlan 901
Switch(config)# monitor session 1 destination interface gigabitethernet0/1
Switch(config)# end
```

Creating an RSPAN Destination Session and Configuring Incoming 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 incoming traffic on the destination port for a network security device (such as a Cisco IDS Sensor Appliance).

For details about the keywords not related to incoming traffic, see the “[Creating an RSPAN Destination Session](#)” section on page 1-20. This procedure assumes that 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.
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 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> untagged vlan <i>vlan-id</i> vlan <i>vlan-id</i> }]}	Specify the SPAN session, the destination port, the packet encapsulation, and the incoming VLAN and encapsulation. For <i>session_number</i> , enter the number defined in Step 4. 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. 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 forwarding of incoming traffic on the destination port and to specify the encapsulation type: <ul style="list-style-type: none"> dot1q vlan <i>vlan-id</i>—Forward incoming packets with IEEE 802.1Q encapsulation with the specified VLAN as the default VLAN. untagged vlan <i>vlan-id</i> or vlan <i>vlan-id</i>—Forward incoming packets with untagged encapsulation type with the specified VLAN as the default VLAN.
Step 5	end	Return to privileged EXEC mode.

	Command	Purpose
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 forwarding of incoming traffic on the interface with VLAN 6 as the default receiving VLAN.

```
Switch(config)# monitor session 2 source remote vlan 901
Switch(config)# monitor session 2 destination interface gigabitethernet0/2 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 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 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.

	Command	Purpose
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 trunk port 2, 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 gigabitethernet0/2 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.



CHAPTER 1

Configuring RMON

This chapter describes how to configure Remote Network Monitoring (RMON) on the Catalyst 2960 and 2960-S switches. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

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

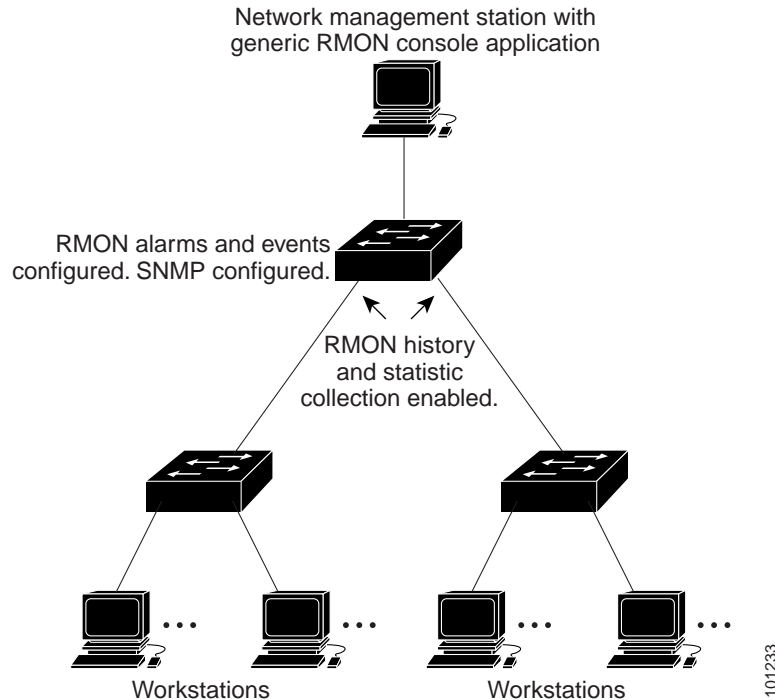
For complete syntax and usage information for the commands used in this chapter, see the “System Management Commands” section in the *Cisco IOS Configuration Fundamentals Command Reference, Release 12.4* on Cisco.com.

- [Understanding RMON, page 1-1](#)
- [Configuring RMON, page 1-2](#)
- [Displaying RMON Status, page 1-6](#)

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 as shown in [Figure 1-1](#).

Figure 1-1 Remote Monitoring Example



The switch supports these RMON groups (defined in RFC 1757):

- Statistics (RMON group 1)—Collects Ethernet statistics (including Fast Ethernet and Gigabit Ethernet statistics, depending on the switch type and supported interfaces) on an interface.
- History (RMON group 2)—Collects a history group of statistics on Ethernet ports (including Fast Ethernet and Gigabit Ethernet statistics, depending on the switch type and supported interfaces) for a specified polling interval.
- Alarm (RMON group 3)—Monitors a specific management information base (MIB) object for a specified interval, triggers an alarm at a specified value (rising threshold), and resets the alarm at another value (falling threshold). Alarms can be used with events; the alarm triggers an event, which can generate a log entry or an SNMP trap.
- Event (RMON group 9)—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.



Note

64-bit counters are not supported for RMON alarms.

Configuring RMON

- [Default RMON Configuration, page 1-3](#)
- [Configuring RMON Alarms and Events, page 1-3](#) (required)

- [Collecting Group History Statistics on an Interface, page 1-5](#) (optional)
- [Collecting Group Ethernet Statistics on an Interface, page 1-5](#) (optional)

Default RMON Configuration

RMON is disabled by default; no alarms or events are configured.

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 the RMON network management capabilities. You must also configure SNMP on the switch to access RMON MIB objects. For more information, see [Chapter 1, “Configuring SNMP.”](#)



Note

64-bit counters are not supported for RMON alarms.

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. <ul style="list-style-type: none"> • For <i>number</i>, specify the alarm number. The range is 1 to 65535. • For <i>variable</i>, specify the MIB object to monitor. • For <i>interval</i>, specify the time in seconds the alarm monitors the MIB variable. The range is 1 to 4294967295 seconds. • Specify the absolute keyword to test each MIB variable directly. Specify the delta keyword to test the change between samples of a MIB variable. • For <i>value</i>, specify a number at which the alarm is triggered and one for when the alarm is reset. The range for the rising threshold and falling threshold 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 string] [log] [owner string] [trap community]	Add an event in the RMON event table that is associated with an RMON event number. <ul style="list-style-type: none"> For <i>number</i>, assign an event number. The range is 1 to 65535. (Optional) For description string, specify a description of the event. (Optional) Use the log keyword to generate an RMON log entry when the event is triggered. (Optional) For owner string, specify the owner of this event. (Optional) For trap community, 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, see 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. <ul style="list-style-type: none"> For <i>index</i>, identify the RMON group of statistics. The range is 1 to 65535. (Optional) For buckets <i>bucket-number</i>, specify the maximum number of buckets desired for the RMON collection history group of statistics. The range is 1 to 65535. The default is 50 buckets. (Optional) For interval <i>seconds</i>, specify the number of seconds in each polling cycle. 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.

	Command	Purpose
Step 3	rmon collection stats <i>index</i> [owner <i>ownername</i>]	Enable RMON statistic collection on the interface. <ul style="list-style-type: none"> For <i>index</i>, specify the RMON group of statistics. The range is from 1 to 65535. (Optional) For owner <i>ownername</i>, enter the name of the owner of the RMON group of statistics.
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 gigabitethernet0/1
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 1-1](#):

Table 1-1 *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, see the “System Management Commands” section in the *Cisco IOS Configuration Fundamentals Command Reference, Release 12.4* on Cisco.com.



CHAPTER 1

Configuring System Message Logging

This chapter describes how to configure system message logging on the Catalyst 2960 and 2960-S switches.



Note

For complete syntax and usage information for the commands used in this chapter, see the *Cisco IOS Configuration Fundamentals Command Reference, Release 12.4*.

- [Understanding System Message Logging, page 1-1](#)
- [Configuring System Message Logging, page 1-2](#)
- [Displaying the Logging Configuration, page 1-14](#)



Caution

Logging messages to the console at a high rate can cause high CPU utilization and adversely affect how the switch operates.

Understanding System Message Logging

By default, a switch sends the output from system messages and **debug** privileged EXEC commands to a logging process. Stack members can trigger system messages. A stack member that generates a system message appends its hostname in the form of *hostname-n*, where *n* is a switch number from 1 to 4, and redirects the output to the logging process on the stack master. Though the stack master is a stack member, it does *not* append its hostname to system messages. 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 appear on the active consoles 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, see 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 a standalone switch, and in the case of a switch stack, on the stack master. If a standalone switch or the stack master fails, the log is lost unless you had saved it to flash memory.

You can access logged system messages by using the switch command-line interface (CLI) or by saving them to a properly configured syslog server. The switch software saves syslog messages in an internal buffer.

You can remotely monitor system messages by viewing the logs on a syslog server or by accessing the switch through Telnet or through the console port. In a switch stack, all stack member consoles provide the same console output.

Configuring System Message Logging

- [System Log Message Format, page 1-2](#)
- [Default System Message Logging Configuration, page 1-4](#)
- [Disabling Message Logging, page 1-4 \(optional\)](#)
- [Setting the Message Display Destination Device, page 1-5 \(optional\)](#)
- [Synchronizing Log Messages, page 1-6 \(optional\)](#)
- [Enabling and Disabling Time Stamps on Log Messages, page 1-8 \(optional\)](#)
- [Enabling and Disabling Sequence Numbers in Log Messages, page 1-8 \(optional\)](#)
- [Defining the Message Severity Level, page 1-9 \(optional\)](#)
- [Limiting Syslog Messages Sent to the History Table and to SNMP, page 1-10 \(optional\)](#)
- [Enabling the Configuration-Change Logger, page 1-11 \(optional\)](#)
- [Configuring UNIX Syslog Servers, page 1-12 \(optional\)](#)

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 appear in this format:

seq no:timestamp: %facility-severity-MNEMONIC:description (hostname-n)

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 1-1 describes the elements of syslog messages.

Table 1-1 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 1-8.
<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 Time Stamps on Log Messages ” section on page 1-8.
<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 1-4 on page 1-14 .
<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 1-3 on page 1-10 .
<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 for a stack master and a stack member (hostname *Switch-2*):

```
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)

00:00:46: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up (Switch-2)
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet2/0/1, changed state to up (Switch-2)
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet2/0/2, changed state to up (Switch-2)
00:00:48: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1, changed state to down
(Switch-2)
00:00:48: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet2/0/1, changed
state to down 2 (Switch-2)
```

This example shows a partial switch system message:

```
00:00:46: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/1, changed state to up
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/2, 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 GigabitEthernet0/1, 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 1-2 Default System Message Logging Configuration

Feature	Default Setting
System message logging to the console	Enabled.
Console severity	Debugging (and numerically lower levels; see Table 1-3 on page 1-10).
Logging file configuration	No filename specified.
Logging buffer size	4096 bytes.
Logging history size	1 message.
Time stamps	Disabled.
Synchronous logging	Disabled.
Logging server	Disabled.
Syslog server IP address	None configured.
Configuration change logger	Disabled
Server facility	Local7 (see Table 1-4 on page 1-14).
Server severity	Informational (and numerically lower levels; see Table 1-3 on page 1-10).

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 console	Disable message logging.
Step 3	end	Return to privileged EXEC mode.

	Command	Purpose
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 appear 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 1-6](#).

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 [<i>size</i>]	Log messages to an internal buffer on a standalone switch or, in the case of a switch stack, on the stack master. The range is 4096 to 2147483647 bytes. The default buffer size is 4096 bytes. If the standalone switch or the stack master fails, 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.
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 1-12 .

	Command	Purpose
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 a standalone switch or, in the case of a switch stack, on the stack master. <ul style="list-style-type: none"> For <i>filename</i>, enter the log message filename. (Optional) For <i>max-file-size</i>, specify the maximum logging file size. The range is 4096 to 2147483647. The default is 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 1-3 on page 1-10. 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.

Use the **logging event power-inline-status** interface configuration command to enable and to disable logging of Power over Ethernet (PoE) events on specific PoE-capable ports. Logging on these ports is enabled by default.

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 is 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 appear, 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. <ul style="list-style-type: none"> 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. <p>You can change the setting of all 16 vty lines at once by entering:</p> <p>line vty 0 15</p> <p>Or you can change the setting of the single vty line being used for your current connection. For example, to change the setting for vty line 2, enter:</p> <p>line vty 2</p> <p>When you enter this command, the mode changes to line configuration.</p>
Step 3	logging synchronous [level <i>severity-level</i> all] limit <i>number-of-buffers</i>]	Enable synchronous logging of messages. <ul style="list-style-type: none"> (Optional) For level severity-level, specify the message severity level. Messages with a severity level equal to or higher than this value are printed asynchronously. Low numbers mean greater severity and high numbers mean lesser severity. The default is 2. (Optional) Specifying level all means that all messages are printed asynchronously regardless of the severity level. (Optional) For limit number-of-buffers, specify the number of buffers to be queued for the terminal after which new messages are dropped. The 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 Time Stamps 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 time stamp 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)
(Switch-2)
```

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 (Switch-2)
```

Enabling and Disabling Sequence Numbers in Log Messages

Because there is a chance that more than one log message can have the same time stamp, you can display messages with sequence numbers so that you can unambiguously see 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.

	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 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 1-3](#).

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 1-3 on page 1-10).
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 1-3 on page 1-10).
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 1-3 on page 1-10). For complete syslog server configuration steps, see the “ Configuring UNIX Syslog Servers ” section on page 1-12.
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 appear 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 1-3 describes the *level* keywords. It also lists the corresponding UNIX syslog definitions from the most severe level to the least severe level.

Table 1-3 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, see 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 1-3 on page 1-10](#)) 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 1-3 on page 1-10 for a list of <i>level</i> keywords. By default, warnings , errors , critical , alerts , and emergencies messages are sent.
Step 3	logging history size number	Specify the number of syslog messages that can be stored in the history table. The default is to store one message. The range is 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 1-3](#) 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.

Enabling the Configuration-Change Logger

You can enable a configuration logger to keep track of configuration changes made with the command-line interface (CLI). When you enter the **logging enable** configuration-change logger configuration command, the log records the session, the user, and the command that was entered to change the configuration. You can configure the size of the configuration log from 1 to 1000 entries (the default is 100). You can clear the log at any time by entering the **no logging enable** command followed by the **logging enable** command to disable and reenabling logging.

Use the **show archive log config {all | number [end-number] | user username [session number] number [end-number] | statistics} [provisioning]** privileged EXEC command to display the complete configuration log or the log for specified parameters.

The default is that configuration logging is disabled.

For information about the commands, see the *Cisco IOS Configuration Fundamentals and Network Management Command Reference, Release 12.3 T*:

http://www.cisco.com/en/US/docs/ios/12_3/configfun/command/reference/cfr_1g04.html

Beginning in privileged EXEC mode, follow these steps to enable configuration logging:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	archive	Enter archive configuration mode.
Step 3	log config	Enter configuration-change logger configuration mode.
Step 4	logging enable	Enable configuration change logging.
Step 5	logging size <i>entries</i>	(Optional) Configure the number of entries retained in the configuration log. The range is from 1 to 1000. The default is 100. Note When the configuration log is full, the oldest log entry is removed each time a new entry is entered.
Step 6	end	Return to privileged EXEC mode.
Step 7	show archive log config	Verify your entries by viewing the configuration log.

This example shows how to enable the configuration-change logger and to set the number of entries in the log to 500.

```
Switch(config)# archive
Switch(config-archive)# log config
Switch(config-archive-log-cfg)# logging enable
Switch(config-archive-log-cfg)# logging size 500
Switch(config-archive-log-cfg)# end
```

This is an example of output for the configuration log:

```
Switch# show archive log config all
  idx  sess      user@line      Logged command
  38   11   unknown user@vty3   |no aaa authorization config-commands
  39   12   unknown user@vty3   |no aaa authorization network default group radius
  40   12   unknown user@vty3   |no aaa accounting dot1x default start-stop group
radius
  41   13   unknown user@vty3   |no aaa accounting system default
  42   14       temi@vty4       |interface GigabitEthernet4/0/1
  43   14       temi@vty4       | switchport mode trunk
  44   14       temi@vty4       | exit
  45   16       temi@vty5       |interface FastEthernet5/0/1
  46   16       temi@vty5       | switchport mode trunk
  47   16       temi@vty5       | exit
```

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 1-4 on page 1-14](#) for information on the facilities. The **debug** keyword specifies the syslog level; see [Table 1-3 on page 1-10](#) 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 1-3 on page 1-10 for <i>level</i> keywords.

	Command	Purpose
Step 4	logging facility <i>facility-type</i>	Configure the syslog facility. See Table 1-4 on page 1-14 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 1-4](#) 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 1-4 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, see the *Cisco IOS Configuration Fundamentals Command Reference, Release 12.4* on Cisco.com.



CHAPTER 1

Configuring SNMP

This chapter describes how to configure the Simple Network Management Protocol (SNMP) on the Catalyst 2960 and 2960-S switches. Unless otherwise noted, the term *switch* refers to a standalone switch and a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.



Note

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release and the *Cisco IOS Network Management Command Reference, Release 12.4*: http://www.cisco.com/en/US/docs/ios/netmgmt/command/reference/nm_book.html

- [Understanding SNMP, page 1-1](#)
- [Configuring SNMP, page 1-6](#)
- [Displaying SNMP Status, page 1-19](#)

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 TCP connection, loss of connection to a neighbor, or other significant events.

The stack master handles the SNMP requests and traps for the whole switch stack. The stack master transparently manages any requests or traps that are related to all stack members. When a new stack master is elected, the new master continues to handle SNMP requests and traps as configured on the previous stack master, assuming that IP connectivity to the SNMP management stations is still in place after the new master has taken control.

For more information about switch stacks, see [Chapter 1, “Managing Switch Stacks.”](#)

These sections contain this conceptual information:

- [SNMP Versions, page 1-2](#)
- [SNMP Manager Functions, page 1-3](#)
- [SNMP Agent Functions, page 1-4](#)
- [SNMP Community Strings, page 1-4](#)
- [Using SNMP to Access MIB Variables, page 1-5](#)
- [SNMP Notifications, page 1-5](#)
- [SNMP ifIndex MIB Object Values, page 1-6](#)

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 cryptographic (encrypted) software image is installed.

Both SNMPv1 and SNMPv2C use a community-based form of security. The community of managers able to access the agent’s MIB is defined by an IP address access control list and password.

SNMPv2C includes a bulk retrieval mechanism and more detailed error message reporting to management stations. The bulk retrieval mechanism retrieves tables and large quantities of information, minimizing the number of round-trips required. The SNMPv2C improved error-handling includes expanded error codes that distinguish different kinds of error conditions; these conditions are reported through a single error code in SNMPv1. Error return codes in SNMPv2C report the error type.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy set up for a user and the group within which the user resides. A security level is the permitted level of security within a security model. A combination of the security level and the security model determine which security mechanism is used when handling an SNMP packet. Available security models are SNMPv1, SNMPv2C, and SNMPv3.

Table 1-1 identifies the characteristics of the different combinations of security models and levels.

Table 1-1 *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	Message Digest 5 (MD5) or Secure Hash Algorithm (SHA)	No	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms.
SNMPv3	authPriv (requires the cryptographic software image)	MD5 or SHA	Data Encryption Standard (DES) or Advanced Encryption Standard (AES)	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms. Allows specifying the User-based Security Model (USM) with these encryption algorithms: <ul style="list-style-type: none"> • DES 56-bit encryption in addition to authentication based on the CBC-DES (DES-56) standard. • 3DES 168-bit encryption • AES 128-bit, 192-bit, or 256-bit encryption

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, SNMPv2C, or SNMPv3.

SNMP Manager Functions

The SNMP manager uses information in the MIB to perform the operations described in Table 1-2.

Table 1-2 *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

Table 1-2 SNMP Operations (continued)

Operation	Description
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
- When a cluster is created, the command switch manages the exchange of messages among member switches and the SNMP application. The Network Assistant software appends the member switch number (*@esN*, where *N* is the switch number) to the first configured RW and RO community strings on the command switch and propagates them to the member switches. For more information, see [Chapter 1, “Clustering Switches”](#) and see *Getting Started with Cisco Network Assistant*, available on Cisco.com.

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 1-1](#), 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.



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.

SNMP ifIndex MIB Object Values

In an NMS, the IF-MIB generates and assigns an interface index (ifIndex) object value that is a unique number greater than zero to identify a physical or a logical interface. When the switch reboots or the switch software is upgraded, the switch uses this same value for the interface. For example, if the switch assigns a port 2 an ifIndex value of 10003, this value is the same after the switch reboots.

The switch uses one of the values in [Table 1-3](#) to assign an ifIndex value to an interface:

Table 1-3 *ifIndex Values*

Interface Type	ifIndex Range
SVI ¹	1–4999
EtherChannel	5001–5048
Physical (such as Gigabit Ethernet or SFP ² -module interfaces) based on type and port numbers	10000–14500
Null	10501 (nonstackable switches) 14501 (stackable switches)
Loopback and Tunnel	24567 +

1. SVI = switch virtual interface

2. SFP = small form-factor pluggable



Note

The switch might not use sequential values within a range.

Configuring SNMP

- [Default SNMP Configuration, page 1-7](#)
- [SNMP Configuration Guidelines, page 1-7](#)
- [Disabling the SNMP Agent, page 1-8](#)
- [Configuring Community Strings, page 1-8](#)
- [Configuring SNMP Groups and Users, page 1-10](#)
- [Configuring SNMP Notifications, page 1-13](#)
- [Setting the CPU Threshold Notification Types and Values, page 1-16](#)
- [Setting the Agent Contact and Location Information, page 1-17](#)
- [Limiting TFTP Servers Used Through SNMP, page 1-17](#)
- [SNMP Examples, page 1-18](#)

Default SNMP Configuration

Table 1-4 shows the default SNMP configuration.

Table 1-4 Default SNMP Configuration

Feature	Default Setting
SNMP agent	Disabled ¹ .
SNMP trap receiver	None configured.
SNMP traps	None enabled except the trap for TCP connections (tty).
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.

1. This is the default when the switch starts and the startup configuration does not have any **snmp-server** global configuration commands.

SNMP Configuration Guidelines

If the switch starts and the switch startup configuration has at least one **snmp-server** global configuration command, the SNMP agent is enabled.

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. See the *Cisco IOS Network Management Command Reference* for information about when you should configure notify views.
- To configure a remote user, specify the IP address or port number for the remote SNMP agent of the device where the user resides.
- Before you configure remote users for a particular agent, configure the SNMP engine ID, using the **snmp-server engineID** global configuration with the **remote** option. The remote agent's SNMP engine ID and user password are used to compute the authentication and privacy digests. If you do not configure the remote engine ID first, the configuration command fails.
- When configuring SNMP informs, you need to configure the SNMP engine ID for the remote agent in the SNMP database before you can send proxy requests or informs to it.
- If a local user is not associated with a remote host, the switch does not send informs for the **auth** (authNoPriv) and the **priv** (authPriv) authentication levels.
- 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 Cisco IOS command exists to enable SNMP. The first **snmp-server** global configuration command that you enter enables all versions of SNMP.

Configuring Community Strings

You use the SNMP community string to define the relationship between the SNMP manager and the agent. The community string acts like a password to permit access to the agent on the switch. Optionally, you can specify one or more of these characteristics associated with the string:

- An access list of IP addresses of the SNMP managers that are permitted to use the community string to gain access to the agent
- A MIB view, which defines the subset of all MIB objects accessible to the given community
- Read and write or read-only permission for the MIB objects accessible to the community

Beginning in privileged EXEC mode, follow these steps to configure a community string on the switch:

	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. Note The @ symbol is used for delimiting the context information. Avoid using the @ symbol as part of the SNMP community string when configuring this command. <ul style="list-style-type: none"> For <i>string</i>, specify a string that acts like a password and permits access to the SNMP protocol. You can configure one or more community strings of any length. (Optional) For view, specify the view record accessible to the community. (Optional) Specify either read-only (ro) if you want authorized management stations to retrieve MIB objects, or specify read-write (rw) if you want authorized management stations to retrieve and modify MIB objects. By default, the community string permits read-only access to all objects. (Optional) For <i>access-list-number</i>, enter an IP standard access list numbered from 1 to 99 and 1300 to 1999.
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. <ul style="list-style-type: none"> For <i>access-list-number</i>, enter the access list number specified in Step 2. The deny keyword denies access if the conditions are matched. The permit keyword permits access if the conditions are matched. For <i>source</i>, enter the IP address of the SNMP managers that are permitted to use the community string to gain access to the agent. (Optional) For <i>source-wildcard</i>, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. <p>Recall that the access list is always terminated by an implicit deny statement for everything.</p>
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. <ul style="list-style-type: none"> The <i>engineid-string</i> is a 24-character ID string with the name of the copy of SNMP. You need not specify the entire 24-character engine ID if it 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 User Datagram Protocol (UDP) port on the remote device. The default is 162.

Command	Purpose
<p>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>]</p>	<p>Configure a new SNMP group on the remote device.</p> <ul style="list-style-type: none"> • For <i>groupname</i>, specify the name of the group. • Specify a security model: <ul style="list-style-type: none"> – v1 is the least secure of the possible security models. – v2c is the second least secure model. It allows transmission of informs and integers twice the normal width. – v3, the most secure, requires you to select an authentication level: <ul style="list-style-type: none"> auth—Enables the Message Digest 5 (MD5) and the Secure Hash Algorithm (SHA) packet authentication. noauth—Enables the noAuthNoPriv security level. This is the default if no keyword is specified. priv—Enables Data Encryption Standard (DES) packet encryption (also called <i>privacy</i>). <p>Note The priv keyword is available only when the cryptographic software image is installed.</p> <ul style="list-style-type: none"> • (Optional) Enter read <i>readview</i> with a string (not to exceed 64 characters) that is the name of the view in which you can only view the contents of the agent. • (Optional) Enter write <i>writeview</i> with a string (not to exceed 64 characters) that is the name of the view in which you enter data and configure the contents of the agent. • (Optional) Enter notify <i>notifyview</i> with a string (not to exceed 64 characters) that is the name of the view in which you specify a notify, inform, or trap. • (Optional) Enter access <i>access-list</i> with a string (not to exceed 64 characters) that is the name of the access list.

	Command	Purpose
Step 4	<code>snmp-server user username groupname {remote host [udp-port port]} {v1 [access access-list] v2c [access access-list] v3 [encrypted] [access access-list] [auth {md5 sha} auth-password]} [priv {des 3des aes {128 192 256}} priv-password]</code>	<p>Add a new user for an SNMP group.</p> <ul style="list-style-type: none"> The <i>username</i> is the name of the user on the host that connects to the agent. The <i>groupname</i> is the name of the group to which the user is associated. Enter remote to specify a remote SNMP entity to which the user belongs and the hostname or IP address of that entity with the optional UDP port number. The default is 162. Enter the SNMP version number (v1, v2c, or v3). If you enter v3, you have these additional options: <ul style="list-style-type: none"> encrypted specifies that the password appears in encrypted format. This keyword is available only when the v3 keyword is specified. auth is an authentication level setting session that can be either the HMAC-MD5-96 (md5) or the HMAC-SHA-96 (sha) authentication level and requires a password string <i>auth-password</i> (not to exceed 64 characters). If you enter v3 and the switch is running the cryptographic software image, you can also configure a private (priv) encryption algorithm and password string <i>priv-password</i> (not to exceed 64 characters). <ul style="list-style-type: none"> priv specifies the User-based Security Model (USM). des specifies the use of the 56-bit DES algorithm. 3des specifies the use of the 168-bit DES algorithm. aes specifies the use of the DES algorithm. You must select either 128-bit, 192-bit, or 256-bit encryption. (Optional) Enter access access-list with a string (not to exceed 64 characters) that is the name of the access list.
Step 5	<code>end</code>	Return to privileged EXEC mode.
Step 6	<code>show running-config</code>	<p>Verify your entries.</p> <p>Note To display SNMPv3 information about auth noauth priv mode configuration, you must enter the show snmp user privileged EXEC command.</p>
Step 7	<code>copy running-config startup-config</code>	(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 Cisco 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 traps, informs, or both. Use the **snmp-server host** global configuration command to specify whether to send SNMP notifications as traps or informs.

Table 1-5 describes the supported switch traps (notification types). You can enable any or all of these traps and configure a trap manager to receive them. To enable the sending of SNMP inform notifications, use the **snmp-server enable traps** global configuration command combined with the **snmp-server host host-addr informs** global configuration command.

Table 1-5 Switch Notification Types

Notification Type Keyword	Description
bridge	Generates STP bridge MIB traps.
cluster	Generates a trap when the cluster configuration changes.
config	Generates a trap for SNMP configuration changes.
copy-config	Generates a trap for SNMP copy configuration changes.
entity	Generates a trap for SNMP entity changes.
cpu threshold	Allow CPU-related traps. This trap is supported only when the switch is running the LAN Base image.
envmon	Generates environmental monitor traps. You can enable any or all of these environmental traps: fan, shutdown, status, supply, temperature.
errdisable	Generates a trap for a port VLAN errdisabled. You can also set a maximum trap rate per minute. The range is from 0 to 10000; the default is 0, which means there is no rate limit.
flash	Generates SNMP FLASH notifications. You can optionally enable notification for flash insertion or removal, which would cause a trap to be issued whenever a switch in the stack is removed or inserted (physical removal, power cycle, or reload).
fru-ctrl	Generates entity field-replaceable unit (FRU) control traps. In the switch stack, this trap refers to the insertion or removal of a switch in the stack.
ipmulticast	Generates a trap for IP multicast routing changes.
mac-notification	Generates a trap for MAC address notifications.
msdp	Generates a trap for Multicast Source Discovery Protocol (MSDP) changes.
ospf	Generates a trap for Open Shortest Path First (OSPF) changes. You can enable any or all of these traps: Cisco specific, errors, link-state advertisement, rate limit, retransmit, and state changes.
pim	Generates a trap for Protocol-Independent Multicast (PIM) changes. You can enable any or all of these traps: invalid PIM messages, neighbor changes, and rendezvous point (RP)-mapping changes.

Table 1-5 Switch Notification Types (continued)

Notification Type Keyword	Description
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. Note When you configure a trap by using the notification type port-security , configure the port security trap first, and then configure the port security trap rate: <ul style="list-style-type: none"> • snmp-server enable traps port-security • snmp-server enable traps port-security trap-rate rate
rtr	Generates a trap for the SNMP Response Time Reporter (RTR). This trap is supported only when the switch is running the LAN Base image.
snmp	Generates a trap for SNMP-type notifications for authentication, cold start, warm start, link up or link down.
storm-control	Generates a trap for SNMP storm-control. You can also set a maximum trap rate per minute. The range is from 0 to 1000; the default is 0 (no limit is imposed; a trap is sent at every occurrence).
stpx	Generates SNMP STP Extended MIB traps.
syslog	Generates SNMP syslog traps.
tty	Generates a trap for TCP connections. This trap is enabled by default.
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 strings, the **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 1-5](#).

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</i> { remote host [udp-port <i>port</i>]} { v1 [access <i>access-list</i>] v2c [access <i>access-list</i>] v3 [encrypted] [access <i>access-list</i>] [auth { md5 sha } <i>auth-password</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. Otherwise, you receive an error message, and the command is not executed.

	Command	Purpose
Step 4	snmp-server group <i>groupname</i> { v1 v2c v3 { auth noauth priv }} [read readview] [write writeview] [notify notifyview] [access access-list]	Configure an SNMP group.
Step 5	snmp-server host <i>host-addr</i> [informs traps] [version { 1 2c 3 { auth noauth priv }}] <i>community-string</i> [<i>notification-type</i>]	<p>Specify the recipient of an SNMP trap operation.</p> <ul style="list-style-type: none"> For <i>host-addr</i>, specify the name or Internet address of the host (the targeted recipient). (Optional) Enter informs to send SNMP informs to the host. (Optional) Enter traps (the default) to send SNMP traps 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. <p>Note The priv keyword is available only when the cryptographic software image is installed.</p> <ul style="list-style-type: none"> For <i>community-string</i>, when version 1 or version 2c is specified, enter the password-like community string sent with the notification operation. When version 3 is specified, enter the SNMPv3 username. <p>Note The @ symbol is used for delimiting the context information. Avoid using the @ symbol as part of the SNMP community string when configuring this command.</p> <ul style="list-style-type: none"> (Optional) For <i>notification-type</i>, use the keywords listed in Table 1-5 on page 1-13. If no type is specified, all notifications are sent.
Step 6	snmp-server enable traps <i>notification-types</i>	<p>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 1-5 on page 1-13, or enter snmp-server enable traps ?</p> <p>To enable multiple types of traps, you must enter a separate snmp-server enable traps command for each trap type.</p> <p>Note When you configure a trap by using the notification type port-security, configure the port security trap first, and then configure the port security trap rate:</p> <ul style="list-style-type: none"> snmp-server enable traps port-security snmp-server enable traps port-security trap-rate <i>rate</i>
Step 7	snmp-server trap-source <i>interface-id</i>	(Optional) Specify the source interface, which provides the IP address for the trap message. This command also sets the source IP address for informs.
Step 8	snmp-server queue-length <i>length</i>	(Optional) Establish the message queue length for each trap host. The range is 1 to 1000; the default is 10.
Step 9	snmp-server trap-timeout <i>seconds</i>	(Optional) Define how often to resend trap messages. The range is 1 to 1000; the default is 30 seconds.
Step 10	end	Return to privileged EXEC mode.

	Command	Purpose
Step 11	show running-config	Verify your entries. Note To display SNMPv3 information about auth noauth priv mode configuration, you must enter the show snmp user privileged EXEC command.
Step 12	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The **snmp-server host** command specifies which hosts receive the notifications. The **snmp-server enable trap** command globally enables the mechanism for the specified notification (for traps and informs). To enable a host to receive an inform, you must configure an **snmp-server host informs** command for the host and globally enable informs by using the **snmp-server enable traps** command.

To remove the specified host from receiving traps, use the **no snmp-server host** *host* global configuration command. The **no snmp-server host** command with no keywords disables traps, but not informs, to the host. To disable informs, use the **no snmp-server host informs** global configuration command. To disable a specific trap type, use the **no snmp-server enable traps** *notification-types* global configuration command.

Setting the CPU Threshold Notification Types and Values

Beginning in privileged EXEC mode, follow these steps to set the CPU threshold notification types and values:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	process cpu threshold type { total process interrupt } rising <i>percentage</i> interval <i>seconds</i> [falling <i>fall-percentage</i> interval <i>seconds</i>]	Set the CPU threshold notification types and values: <ul style="list-style-type: none"> • total—set the notification type to total CPU utilization. • process—set the notification type to CPU process utilization. • interrupt—set the notification type to CPU interrupt utilization. • rising <i>percentage</i>—the percentage (1 to 100) of CPU resources that, when exceeded for the configured interval, sends a CPU threshold notification. • interval <i>seconds</i>—the duration of the CPU threshold violation in seconds (5 to 86400) that, when met, sends a CPU threshold notification. • falling <i>fall-percentage</i>—the percentage (1 to 100) of CPU resources that, when usage falls below this level for the configured interval, sends a CPU threshold notification. <p>This value must be equal to or less than the rising <i>percentage</i> value. If not specified, the falling <i>fall-percentage</i> value is the same as the rising <i>percentage</i> value.</p>
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.

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: <code>snmp-server contact Dial System Operator at beeper 21555.</code>
Step 3	snmp-server location <i>text</i>	Set the system location string. For example: <code>snmp-server location Building 3/Room 222</code>
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.
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. <ul style="list-style-type: none"> For <i>access-list-number</i>, enter the access list number specified in Step 2. The deny keyword denies access if the conditions are matched. The permit keyword permits access if the conditions are matched. For <i>source</i>, enter the IP address of the TFTP servers that can access the switch. (Optional) For <i>source-wildcard</i>, enter the wildcard bits, in dotted decimal notation, to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.

	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.

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
```

This example shows how to associate a user with a remote host and to send **auth** (authNoPriv) authentication-level informs when the user enters global configuration mode:

```
Switch(config)# snmp-server engineID remote 192.180.1.27 00000063000100a1c0b4011b
Switch(config)# snmp-server group authgroup v3 auth
Switch(config)# snmp-server user authuser authgroup remote 192.180.1.27 v3 auth md5
mypassword
Switch(config)# snmp-server user authuser authgroup v3 auth md5 mypassword
Switch(config)# snmp-server host 192.180.1.27 informs version 3 auth authuser config
Switch(config)# snmp-server enable traps
Switch(config)# snmp-server inform retries 0
```

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 also can use the other privileged EXEC commands in [Table 1-6](#) to display SNMP information. For information about the fields in the displays, see the *Cisco IOS Configuration Fundamentals Command Reference*.

Table 1-6 *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 pending	Displays information on pending SNMP requests.
show snmp sessions	Displays information on the current SNMP sessions.
show snmp user	Displays information on each SNMP user name in the SNMP users table. Note You must use this command to display SNMPv3 configuration information for auth noauth priv mode. This information is not displayed in the show running-config output.



CHAPTER 1

Configuring Cisco IOS IP SLAs Operations

This chapter describes how to use Cisco IOS IP Service Level Agreements (SLAs) on the Catalyst 2960 and 2960-S switches. Cisco IP SLAs is a part of Cisco IOS software that allows Cisco customers to analyze IP service levels for IP applications and services by using active traffic monitoring—the generation of traffic in a continuous, reliable, and predictable manner—for measuring network performance. With Cisco IOS IP SLAs, service provider customers can measure and provide service level agreements, and enterprise customers can verify service levels, verify outsourced service level agreements, and understand network performance. Cisco IOS IP SLAs can perform network assessments, verify quality of service (QoS), ease the deployment of new services, and assist with network troubleshooting. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note

To use Cisco IOS IP Service Level Agreements (SLAs), the switch must be running the LAN Base image.

The switch supports only IP SLAs responder functionality and must be configured with another device that supports full IP SLAs functionality.

For more information about IP SLAs, see the *Cisco IOS IP SLAs Configuration Guide, Release 12.4T*: http://www.cisco.com/en/US/docs/ios/ipsla/configuration/guide/12_4t/sla_12_4t_book.html

For command syntax information, see the command reference: http://www.cisco.com/en/US/docs/ios/ipsla/command/reference/sla_book.html

This chapter has these sections:

- [Understanding Cisco IOS IP SLAs, page 1-1](#)
- [Configuring IP SLAs Operations, page 1-5](#)
- [Monitoring IP SLAs Operations, page 1-6](#)

Understanding Cisco IOS IP SLAs

Cisco IOS IP SLAs sends data across the network to measure performance between multiple network locations or across multiple network paths. It simulates network data and IP services and collects network performance information in real time. Cisco IOS IP SLAs generates and analyzes traffic either between Cisco IOS devices or from a Cisco IOS device to a remote IP device such as a network application server. Measurements provided by the various Cisco IOS IP SLAs operations can be used for troubleshooting, for problem analysis, and for designing network topologies.

Depending on the specific Cisco IOS IP SLAs operation, various network performance statistics are monitored within the Cisco device and stored in both command-line interface (CLI) and Simple Network Management Protocol (SNMP) MIBs. IP SLAs packets have configurable IP and application layer options such as source and destination IP address, User Datagram Protocol (UDP)/TCP port numbers, a type of service (ToS) byte (including Differentiated Services Code Point [DSCP] and IP Prefix bits), Virtual Private Network (VPN) routing/forwarding instance (VRF), and URL web address.

Because Cisco IP SLAs is Layer 2 transport independent, you can configure end-to-end operations over disparate networks to best reflect the metrics that an end user is likely to experience. IP SLAs collects a unique subset of these performance metrics:

- Delay (both round-trip and one-way)
- Jitter (directional)
- Packet loss (directional)
- Packet sequencing (packet ordering)
- Path (per hop)
- Connectivity (directional)
- Server or website download time

Because Cisco IOS IP SLAs is SNMP-accessible, it can also be used by performance-monitoring applications like CiscoWorks Internetwork Performance Monitor (IPM) and other third-party Cisco partner performance management products. You can find more details about network management products that use Cisco IOS IP SLAs:

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

Using IP SLAs can provide these benefits:

- Service-level agreement monitoring, measurement, and verification.
- Network performance monitoring
 - Measures the jitter, latency, or packet loss in the network.
 - Provides continuous, reliable, and predictable measurements.
- IP service network health assessment to verify that the existing QoS is sufficient for new IP services.
- Edge-to-edge network availability monitoring for proactive verification and connectivity testing of network resources (for example, shows the network availability of an NFS server used to store business critical data from a remote site).
- Troubleshooting of network operation by providing consistent, reliable measurement that immediately identifies problems and saves troubleshooting time.
- Multiprotocol Label Switching (MPLS) performance monitoring and network verification (if the switch supports MPLS)

This section has this information about IP SLAs functionality:

- [Using Cisco IOS IP SLAs to Measure Network Performance, page 1-3](#)
- [IP SLAs Responder and IP SLAs Control Protocol, page 1-4](#)
- [Response Time Computation for IP SLAs, page 1-4](#)

Using Cisco IOS IP SLAs to Measure Network Performance

You can use IP SLAs to monitor the performance between any area in the network—core, distribution, and edge—without deploying a physical probe. It uses generated traffic to measure network performance between two networking devices. Figure 1-1 shows how IP SLAs begins when the source device sends a generated packet to the destination device. After the destination device receives the packet, depending on the type of IP SLAs operation, it responds with time-stamp information for the source to make the calculation on performance metrics. An IP SLAs operation performs a network measurement from the source device to a destination in the network using a specific protocol such as UDP.

Figure 1-1 Cisco IOS IP SLAs Operation

To implement IP SLAs network performance measurement, you need to perform these tasks:

1. Enable the IP SLAs responder, if required.
2. Configure the required IP SLAs operation type.
3. Configure any options available for the specified operation type.
4. Configure threshold conditions, if required.
5. Schedule the operation to run, then let the operation run for a period of time to gather statistics.
6. Display and interpret the results of the operation using the Cisco IOS CLI or a network management system (NMS) system with SNMP.

For more information about IP SLAs operations, see the operation-specific chapters in the *Cisco IOS IP SLAs Configuration Guide*:

http://www.cisco.com/en/US/docs/ios/ipsla/configuration/guide/12_4t/sla_12_4t_book.html



Note

The switch does not support Voice over IP (VoIP) service levels using the gatekeeper registration delay operations measurements. Before configuring any IP SLAs application, you can use the **show ip sla application** privileged EXEC command to verify that the operation type is supported on your software image.

IP SLAs Responder and IP SLAs Control Protocol

The IP SLAs responder is a component embedded in the destination Cisco device that allows the system to anticipate and respond to IP SLAs request packets. The responder provides accurate measurements without the need for dedicated probes. The responder uses the Cisco IOS IP SLAs Control Protocol to provide a mechanism through which it can be notified on which port it should listen and respond. Only a Cisco IOS device can be a source for a destination IP SLAs Responder.



Note

The IP SLAs responder can be a Cisco IOS Layer 2, responder-configurable switch, such as a Catalyst 2960 or IE 3000 switch running the LAN base image, or a Catalyst 3560 or 3750 switch running the IP base image. The responder does not need to support full IP SLAs functionality.

Figure 1-1 shows where the Cisco IOS IP SLAs responder fits in the IP network. The responder listens on a specific port for control protocol messages sent by an IP SLAs operation. Upon receipt of the control message, it enables the specified UDP or TCP port for the specified duration. During this time, the responder accepts the requests and responds to them. It disables the port after it responds to the IP SLAs packet, or when the specified time expires. MD5 authentication for control messages is available for added security.

You do not need to enable the responder on the destination device for all IP SLAs operations. For example, a responder is not required for services that are already provided by the destination router (such as Telnet or HTTP). You cannot configure the IP SLAs responder on non-Cisco devices and Cisco IOS IP SLAs can send operational packets only to services native to those devices.


Response Time Computation for IP SLAs

Switches and routers can take tens of milliseconds to process incoming packets due to other high priority processes. This delay affects the response times because the test-packet reply might be in a queue while waiting to be processed. In this situation, the response times would not accurately represent true network delays. IP SLAs minimizes these processing delays on the source device as well as on the target device (if the responder is being used) to determine true round-trip times. IP SLAs test packets use time stamping to minimize the processing delays.

When the IP SLAs responder is enabled, it allows the target device to take time stamps when the packet arrives on the interface at interrupt level and again just as it is leaving, eliminating the processing time. This time stamping is made with a granularity of sub-milliseconds (ms).

Figure 1-2 demonstrates how the responder works. Four time stamps are taken to make the calculation for round-trip time. At the target router, with the responder functionality enabled, time stamp 2 (TS2) is subtracted from time stamp 3 (TS3) to produce the time spent processing the test packet as represented by delta. This delta value is then subtracted from the overall round-trip time. Notice that the same principle is applied by IP SLAs on the source router where the incoming time stamp 4 (TS4) is also taken at the interrupt level to allow for greater accuracy.

Figure 1-2 Cisco IOS IP SLAs Responder Time Stamping



An additional benefit of the two time stamps at the target device is the ability to track one-way delay, jitter, and directional packet loss. Because much network behavior is asynchronous, it is critical to have these statistics. However, to capture one-way delay measurements, you must configure both the source router and target router with Network Time Protocol (NTP) so that the source and target are synchronized to the same clock source. One-way jitter measurements do not require clock synchronization.

Configuring IP SLAs Operations

This section does not include configuration information for all available operations as the configuration information details are included in the *Cisco IOS IP SLAs Configuration Guide*. It includes only the procedure for configuring the responder, as the switch includes only responder support.

For details about configuring other operations, see the *Cisco IOS IP SLAs Configuration Guide*: http://www.cisco.com/en/US/docs/ios/ipsla/configuration/guide/12_4t/sla_12_4t_book.html

This section includes this information:

- [Default Configuration, page 1-5](#)
- [Configuration Guidelines, page 1-5](#)
- [Configuring the IP SLAs Responder, page 1-6](#)

Default Configuration

No IP SLAs operations are configured.

Configuration Guidelines

For information on the IP SLAs commands, see the *Cisco IOS IP SLAs Command Reference, Release 12.4T* command reference:

http://www.cisco.com/en/US/docs/ios/ipsla/command/reference/sla_book.html

For detailed descriptions and configuration procedures, see the *Cisco IOS IP SLAs Configuration Guide, Release 12.4T*:

http://www.cisco.com/en/US/docs/ios/ipsla/configuration/guide/12_4t/sla_12_4t_book.html

Configuring the IP SLAs Responder

The IP SLAs responder is available only on Cisco IOS software-based devices, including some Layer 2 switches that do not support full IP SLAs functionality, such as the Catalyst 2960 or the Cisco ME 2400 or IE 3000 switch running the LAN base image. Beginning in privileged EXEC mode, follow these steps to configure the IP SLAs responder on the target device (the operational target):

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip sla responder { tcp-connect udp-echo } ipaddress <i>ip-address</i> port <i>port-number</i>	Configure the switch as an IP SLAs responder. The optional keywords have these meanings: <ul style="list-style-type: none"> • tcp-connect—Enable the responder for TCP connect operations. • udp-echo—Enable the responder for User Datagram Protocol (UDP) echo or jitter operations. • ipaddress <i>ip-address</i>—Enter the destination IP address. • port <i>port-number</i>—Enter the destination port number. <p>Note The IP address and port number must match those configured on the source device for the IP SLAs operation.</p>
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip sla responder	Verify the IP SLAs responder configuration on the device.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the IP SLAs responder, enter the **no ip sla responder** global configuration command. This example shows how to configure the device as a responder for the UDP jitter IP SLAs operation in the next procedure:

```
Switch(config)# ip sla responder udp-echo 172.29.139.134 5000
```



Note

For the IP SLAs responder to function, you must also configure a source device, such as a Catalyst 3750 or Catalyst 3560 switch running the IP services image, that has full IP SLAs support. Refer to the documentation for the source device for configuration information.

Monitoring IP SLAs Operations

Use the User EXEC or Privileged EXEC commands in [Table 1-1](#) to display IP SLAs operations configuration.

Table 1-1 Monitoring IP SLAs Operations

Command	Purpose
show ip sla authentication	Display IP SLAs authentication information.
show ip sla responder	Display information about the IP SLAs responder.



CHAPTER 1

Configuring Network Security with ACLs

This chapter describes how to configure network security on the Catalyst 2960 and 2960-S switches by using access control lists (ACLs), also referred to as access lists. Unless otherwise noted, the term *switch* refers to a standalone switch and a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.



Note

If the switch is running the LAN Lite image, you can configure ACLs, but you cannot attach them to physical interfaces. When it is running either the LAN Lite or LAN base image, you can attach ACLs to VLAN interfaces to filter traffic to the CPU.

In this chapter, references to IP ACLs are specific to IP Version 4 (IPv4) ACLs.

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release, the “Configuring IP Services” section in the “IP Addressing and Services” chapter of the *Cisco IOS IP Configuration Guide, Release 12.4*, and the *Cisco IOS IP Command Reference, Volume 1 of 3: Addressing and Services, Release 12.4* on Cisco.com.

This chapter includes these sections:

- [Understanding ACLs, page 1-1](#)
- [Configuring IPv4 ACLs, page 1-6](#)
- [Creating Named MAC Extended ACLs, page 1-22](#)
- [Displaying IPv4 ACL Configuration, page 1-24](#)

Understanding ACLs

Packet filtering can help limit network traffic and restrict network use by certain users or devices. ACLs filter traffic as it passes through a switch and permit or deny packets crossing specified interfaces. An ACL is a sequential collection of permit and deny conditions that apply to packets. When a packet is received on an interface, the switch compares the fields in the packet against any applied ACLs to verify that the packet has the required permissions to be forwarded, based on the criteria specified in the access lists. One by one, it tests packets against the conditions in an access list. The first match decides whether the switch accepts or rejects the packets. Because the switch stops testing after the first match, the order

of conditions in the list is critical. If no conditions match, the switch rejects the packet. If there are no restrictions, the switch forwards the packet; otherwise, the switch drops the packet. The switch can use ACLs on all packets it forwards.

You configure access lists on a 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. For example, you can allow e-mail traffic to be forwarded but not Telnet traffic.

An ACL contains an ordered list of access control entries (ACEs). Each ACE specifies *permit* or *deny* and a set of conditions the packet must satisfy in order to match the ACE. The meaning of *permit* or *deny* depends on the context in which the ACL is used.

The switch supports IP ACLs and Ethernet (MAC) ACLs:

- IP ACLs filter IPv4 traffic, including TCP, User Datagram Protocol (UDP), Internet Group Management Protocol (IGMP), and Internet Control Message Protocol (ICMP).
- Ethernet ACLs filter non-IP traffic.



Note MAC ACLs are supported only when the switch is running the LAN base image.

This switch also supports quality of service (QoS) classification ACLs. For more information, see the [“Classification Based on QoS ACLs”](#) section on page 1-8.

These sections contain this conceptual information:

- [Supported ACLs](#), page 1-2
- [Handling Fragmented and Unfragmented Traffic](#), page 1-4
- [ACLs and Switch Stacks](#), page 1-5

Supported ACLs

- 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. For more information, see the [“Port ACLs”](#) section on page 1-3.
- Router ACLs access-control routed traffic between VLANs and are applied to Layer 3 interfaces in a specific direction (inbound or outbound). For more information, see the [“Router ACLs”](#) section on page 1-4.



Note Router ACLs are supported only on SVIs.

You can use input port ACLs and router ACLs on the same switch. However, a port ACL takes precedence over a router ACL.

- When an input router ACL and input port ACL exist in a 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.

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 and can be applied only on interfaces in the inbound direction. These access lists are supported:

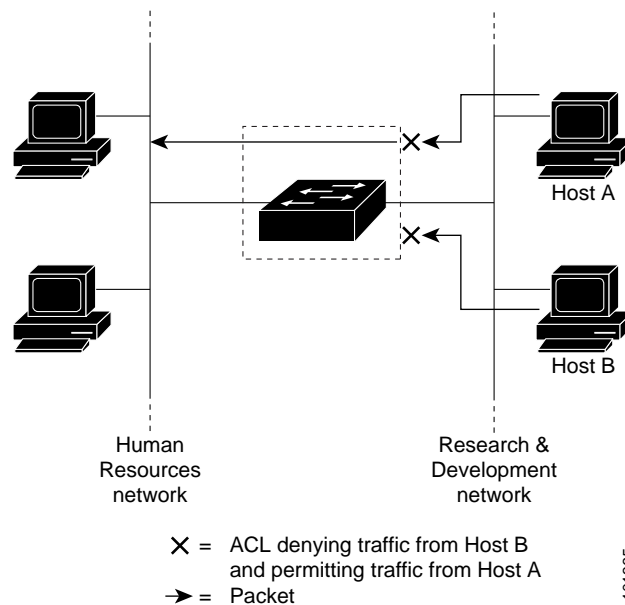
- Standard IP access lists using source addresses
- Extended IP access lists using source and destination addresses and optional protocol type information
- MAC extended access lists using source and destination MAC addresses and optional protocol type information



Note MAC ACLs are supported only when the switch is running the LAN base image.

The switch examines ACLs associated with all inbound features configured on a given interface and permits or denies packet forwarding based on how the packet matches the entries in the ACL. In this way, ACLs control access to a network or to part of a network. [Figure 1-1](#) is an example of using port ACLs to control access to a network when all workstations are 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. Port ACLs can only be applied to Layer 2 interfaces in the inbound direction.

Figure 1-1 Using ACLs to Control Traffic to a 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 IP traffic by using IP access lists and non-IP 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.

Router ACLs

You can apply router ACLs on switch virtual interfaces (SVIs), which are Layer 3 interfaces to VLANs. You apply router ACLs on interfaces for specific directions (inbound or outbound). You can apply one router ACL in each direction on an interface.

An 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.

Supported access lists for IPv4 traffic:

- Standard IP access lists use source addresses for matching operations.
- Extended IP access lists use source and destination addresses and optional protocol information for matching operations.

As with port ACLs, the switch examines ACLs associated with features configured on a given interface. However, you can apply only inbound port ACLs, while router ACLs are supported in both directions. 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 1-1](#), ACLs applied at the router input allow Host A to access the Human Resources network but prevent Host B from accessing the same network.

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.

ACLs and Switch Stacks

**Note**

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

ACL support is the same for a switch stack as for a standalone switch. ACL configuration information is propagated to all switches in the stack. All switches in the stack, including the stack master, process the information and program their hardware. (For information about switch stacks, see [Chapter 1, “Managing Switch Stacks.”](#))

The stack master performs these ACL functions:

- It processes the ACL configuration and propagates the information to all stack members.
- It distributes the ACL information to any switch that joins the stack.
- If packets must be forwarded by software for any reason (for example, not enough hardware resources), the master switch forwards the packets only after applying ACLs on the packets.
- It programs its hardware with the ACL information it processes.

Stack members perform these ACL functions:

- They receive the ACL information from the master switch and program their hardware.
- They act as standby switches, ready to take over the role of the stack master if the existing master were to fail and they were to be elected as the new stack master.

When a stack master fails and a new stack master is elected, the newly elected master reparses the backed up running configuration. (See [Chapter 1, “Managing Switch Stacks.”](#)) The ACL configuration that is part of the running configuration is also reparsed during this step. The new stack master distributes the ACL information to all switches in the stack.

Configuring IPv4 ACLs



Note

If the switch is running the LAN Lite image, you can configure ACLs, but you cannot attach them to physical interfaces. When running either the LAN Lite or LAN base image, you can attach ACLs to VLAN interfaces to filter traffic to the CPU.

Configuring IP v4ACLs on the switch is the same as configuring IPv4 ACLs on other Cisco switches and routers. The process is briefly described here. For more detailed information on configuring ACLs, see the “Configuring IP Services” section in the “IP Addressing and Services” chapter of the *Cisco IOS IP Configuration Guide, Release 12.4*. For detailed information about the commands, see the *Cisco IOS IP Command Reference, Volume 1 of 3: Addressing and Services, Release 12.4* on Cisco.com.

The switch does not support these Cisco IOS router ACL-related features:

- Non-IP protocol ACLs (see [Table 1-1 on page 1-7](#)) 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

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 the access conditions.
- Step 2** Apply the ACL to interfaces or terminal lines.
-

These sections contain this configuration information:

- [Creating Standard and Extended IPv4 ACLs, page 1-7](#)
- [Applying an IPv4 ACL to a Terminal Line, page 1-17](#)
- [Applying an IPv4 ACL to an Interface, page 1-17](#)
- [Hardware and Software Treatment of IP ACLs, page 1-19](#)
- [Troubleshooting ACLs, page 1-19](#)
- [IPv4 ACL Configuration Examples, page 1-20](#)

Creating Standard and Extended IPv4 ACLs

This section describes IP ACLs. An ACL is a sequential collection of permit and deny conditions. One by one, the switch tests packets against the conditions in an access list. The first match determines whether the switch accepts or rejects the packet. Because the switch stops testing 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 IPv4:

- 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 1-7](#)
- [Creating a Numbered Standard ACL, page 1-8](#)
- [Creating a Numbered Extended ACL, page 1-9](#)
- [Resequencing ACEs in an ACL, page 1-13](#)
- [Creating Named Standard and Extended ACLs, page 1-13](#)
- [Using Time Ranges with ACLs, page 1-15](#)
- [Including Comments in ACLs, page 1-16](#)

Access List Numbers

The number you use to denote your ACL shows the type of access list that you are creating. [Table 1-1](#) lists the access-list number and corresponding access list type and shows whether or not they are supported in the switch. The switch supports IPv4 standard and extended access lists, numbers 1 to 199 and 1300 to 2699.

Table 1-1 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

Table 1-1 Access List Numbers (continued)

Access List Number	Type	Supported
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 by using the supported numbers. That is, the name of a standard IP ACL can be 1 to 99; the name of an extended IP ACL can be 100 to 199. The advantage of using named ACLs instead of numbered lists is that you can delete individual entries from a named list.

Creating a Numbered Standard ACL

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>]	<p>Define a standard IPv4 access list by using a source address and wildcard.</p> <p>The <i>access-list-number</i> is a decimal number from 1 to 99 or 1300 to 1999.</p> <p>Enter deny or permit to specify whether to deny or permit access if conditions are matched.</p> <p>The <i>source</i> is the source address of the network or host from which the packet is being sent specified as:</p> <ul style="list-style-type: none"> The 32-bit quantity in dotted-decimal format. The keyword any as an abbreviation for <i>source</i> and <i>source-wildcard</i> of 0.0.0.0 255.255.255.255. You do not need to enter a source-wildcard. The keyword host as an abbreviation for <i>source</i> and <i>source-wildcard</i> of <i>source</i> 0.0.0.0. <p>Note (Optional) The <i>source-wildcard</i> applies wildcard bits to the source.</p>
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
 10 deny 171.69.198.102
 20 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.

After creating a numbered standard IPv4 ACL, you can apply it to terminal lines (see the “[Applying an IPv4 ACL to a Terminal Line](#)” section on page 1-17) and to interfaces (see the “[Applying an IPv4 ACL to an Interface](#)” section on page 1-17).

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 (**igrp**), Encapsulation Security Payload (**esp**), generic routing encapsulation (**gre**), Internet Control Message Protocol (**icmp**), Internet Group Management Protocol (**igmp**), 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 for each protocol, see these command references:

- *Cisco IOS IP Command Reference, Volume 1 of 3: Addressing and Services, Release 12.4*
- *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols, Release 12.4*
- *Cisco IOS IP Command Reference, Volume 3 of 3: Multicast, Release 12.4*

**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] [time-range <i>time-range-name</i>] [dscp <i>dscp</i>]	<p>Define an extended IPv4 access list and the access conditions.</p> <p>The <i>access-list-number</i> is a decimal number from 100 to 199 or 2000 to 2699.</p> <p>Enter deny or permit to specify whether to deny or permit the packet if conditions are matched.</p> <p>For <i>protocol</i>, enter the name or number of an IP protocol: 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.</p> <p>Note This step includes options for most IP protocols. For additional specific parameters for TCP, UDP, ICMP, and IGMP, see steps 2b through 2e.</p> <p>The <i>source</i> is the number of the network or host from which the packet is sent.</p> <p>The <i>source-wildcard</i> applies wildcard bits to the source.</p> <p>The <i>destination</i> is the network or host number to which the packet is sent.</p> <p>The <i>destination-wildcard</i> applies wildcard bits to the destination.</p> <p>Source, source-wildcard, destination, and destination-wildcard can be specified as:</p> <ul style="list-style-type: none"> • 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. <p>The other keywords are optional and have these meanings:</p> <ul style="list-style-type: none"> • 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). • time-range—For an explanation of this keyword, see the “Using Time Ranges with ACLs” section on page 1-15. • 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.
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 .	

	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] [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] [time-range <i>time-range-name</i>] [dscp <i>dscp</i>]	Define an extended IP access list by using an abbreviation for a source and a 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 the 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] [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 source-wildcard</i>) or destination (if positioned after <i>destination 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 see the “Configuring IP Services” section in the “IP Addressing and Services” chapter of the <i>Cisco IOS IP Configuration Guide, Release 12.4</i> . Use only TCP port numbers or names when filtering TCP. The other optional keywords have these meanings: <ul style="list-style-type: none"> • 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] [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 the [<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	<code>access-list access-list-number {deny permit} icmp source source-wildcard destination destination-wildcard [icmp-type [[icmp-type icmp-code] [icmp-message]] [precedence precedence] [tos tos] [fragments] [time-range time-range-name] [dscp dscp]</code>	<p>(Optional) Define an extended ICMP access list and the access conditions.</p> <p>Enter icmp for Internet Control Message Protocol.</p> <p>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.</p> <p>These optional keywords have these meanings:</p> <ul style="list-style-type: none"> <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 the ICMP message code type, a number from 0 to 255. <i>icmp-message</i>—Enter to filter ICMP packets by the ICMP message type name or the ICMP message type and code name. To see a list of ICMP message type names and code names, use the <code>?</code>, or see the “Configuring IP Services” section of the <i>Cisco IOS IP Configuration Guide, Release 12.4</i>.
Step 2e	<code>access-list access-list-number {deny permit} igmp source source-wildcard destination destination-wildcard [igmp-type] [precedence precedence] [tos tos] [fragments] [time-range time-range-name] [dscp dscp]</code>	<p>(Optional) Define an extended IGMP access list and the access conditions.</p> <p>Enter igmp for Internet Group Management Protocol.</p> <p>The IGMP parameters are the same as those described for most IP protocols in Step 2a, with this optional parameter.</p> <p><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).</p>
Step 3	<code>end</code>	Return to privileged EXEC mode.
Step 4	<code>show access-lists [number name]</code>	Verify the access list configuration.
Step 5	<code>copy running-config startup-config</code>	(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 to 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
 10 deny tcp 171.69.198.0 0.0.0.255 172.20.52.0 0.0.0.255 eq telnet
 20 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 IPv4 ACL to a Terminal Line” section on page 1-17), to interfaces (see the “Applying an IPv4 ACL to an Interface” section on page 1-17).

Resequencing ACEs in an ACL

Sequence numbers for the entries in an access list are automatically generated when you create a new ACL. You can use the **ip access-list resequence** global configuration command to edit the sequence numbers in an ACL and change the order in which ACEs are applied. For example, if you add a new ACE to an ACL, it is placed at the bottom of the list. By changing the sequence number, you can move the ACE to a different position in the ACL.

For information about the **ip access-list resequence** command:

http://www.cisco.com/en/US/docs/ios/12_2s/feature/guide/fsaclseq.html#wp1027188

Creating Named Standard and Extended ACLs

You can identify IPv4 ACLs with an alphanumeric string (a name) rather than a number. You can use named ACLs to configure more IPv4 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.
- 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 IPv4 ACLs” section on page 1-7.

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 IPv4 access list using a name, and enter access-list configuration mode. 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 } or permit { <i>source</i> [<i>source-wildcard</i>] host <i>source</i> any }	In access-list configuration mode, specify one or more conditions denied or permitted to decide if the packet is forwarded or dropped. <ul style="list-style-type: none"> • 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.

	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.

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 IPv4 access list using a name, and enter access-list configuration mode. The name can be a number from 100 to 199.
Step 3	{deny permit} <i>protocol</i> { <i>source</i> [<i>source-wildcard</i>] host <i>source</i> any } { <i>destination</i> [<i>destination-wildcard</i>] host <i>destination</i> any } [precedence <i>precedence</i>] [tos <i>tos</i>] [established] [time-range <i>time-range-name</i>]	In access-list configuration mode, specify the conditions allowed or denied. See the “ Creating a Numbered Extended ACL ” section on page 1-9 for definitions of protocols and other keywords. <ul style="list-style-type: none"> host <i>source</i>—A source and source wildcard of <i>source</i> 0.0.0.0. host <i>destination</i>—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 IPv4 ACL to an Interface](#)” section on page 1-17).

Using Time Ranges with ACLs

You can selectively apply extended ACLs based on the time of day and the 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 IPv4 ACLs](#)” section on page 1-7, and the “[Creating Named Standard and Extended ACLs](#)” section on page 1-13.

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 1-2.

Beginning in privileged EXEC mode, follow these steps to configure a 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 [start <i>time date</i>] [end <i>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. <ul style="list-style-type: none"> 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. See 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 to configure January 1, 2006, as a company holiday 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_2006
Switch(config-time-range)# absolute start 00:00 1 Jan 2006 end 23:59 1 Jan 2006
Switch(config-time-range)# end
Switch# show time-range
time-range entry: new_year_day_2003 (inactive)
    absolute start 00:00 01 January 2006 end 23:59 01 January 2006
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_2006
Switch(config)# access-list 188 permit tcp any any time-range workhours
Switch(config)# end
Switch# show access-lists
Extended IP access list 188
  10 deny tcp any any time-range new_year_day_2006 (inactive)
  20 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_2006
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 lpip_default
  10 permit ip any any
Extended IP access list deny_access
  10 deny tcp any any time-range new_year_day_2006 (inactive)
Extended IP access list may_access
  10 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 that belongs to Jones is allowed access, and the workstation that belongs 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 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 IPv4 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 IPv4 ACL to an Interface” section on page 1-17](#).

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. <ul style="list-style-type: none"> console—Specify the console terminal line. The console port is DCE. vty—Specify a virtual terminal for remote console access. <p>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.</p>
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.
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 IPv4 ACL to an Interface

Note these guidelines:

- Apply an ACL only to inbound Layer 2 ports.
- Apply an ACL to either inbound or outbound VLAN interfaces to filter packets that are intended for the CPU, such as SNMP, Telnet, or web traffic. IPv4 ACLs applied to VLAN interfaces provide switch management security by limiting access to a specific host in the network or to specific applications (SNMP, Telnet, SSH, and so on). ACLs attached to VLAN interfaces do not impact the hardware switching of packets on the VLAN.



Note On switches running the LAN Lite image, you can apply ACLs only to VLAN interfaces and not to physical interfaces.

- Apply an ACL to either outbound or inbound Layer 3 SVIs.
- When controlling access to an interface, you can use a named or numbered ACL.
- If you apply an ACL to a port that is a member of a VLAN, the port ACL takes precedence over an ACL applied to the VLAN interface.
- 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. The port ACL always filters incoming packets received on the Layer 2 port.
- If you apply an ACL to a Layer 3 interface and routing is not enabled, 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.
- When you configure an egress ACL to permit traffic with a particular DSCP value, you must use the original DSCP value instead of a rewritten value.

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 <i>interface-id</i>	Identify a specific interface for configuration, and enter interface configuration mode. On switches running the LAN base image, the interface can be a physical interface or VLAN interface. On switches running the LAN Lite image, the interface must be a VLAN interface.
Step 3	ip access-group { <i>access-list-number</i> <i>name</i> } { in out }	Control access to the specified interface. The out keyword is supported only for VLAN interfaces.
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 to a port to filter packets entering the port:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip access-group 2 in
```

This example shows how to apply access list 3 to filter packets going to the CPU:

```
Switch(config)# interface vlan 1
Switch(config-if)# ip access-group 3 in
```



Note When you apply the **ip access-group** interface configuration command to a Layer 3 SVI, the interface must have an IP address. Layer 3 access groups filter packets that are routed or are received by Layer 3 processes on the CPU.

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 sending 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.

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.

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 packets.

Troubleshooting ACLs

If this ACL manager message appears and [chars] is the access-list name,

```
ACLMGR-2-NOVMR: Cannot generate hardware representation of access list [chars]
```

The switch has insufficient resources to create a hardware representation of the ACL. The resources include hardware memory and label space but not CPU memory. A lack of available logical operation units or specialized hardware resources causes this problem. Logical operation units are needed for a TCP flag match or a test other than **eq** (**ne**, **gt**, **lt**, or **range**) on TCP, UDP, or SCTP port numbers.

Use one of these workarounds:

- Modify the ACL configuration to use fewer resources.
- Rename the ACL with a name or number that alphanumerically precedes the ACL names or numbers.

To determine the specialized hardware resources, enter the **show platform layer4 acl map** privileged EXEC command. If the switch does not have available resources, the output shows that index 0 to index 15 are not available.

For more information about configuring ACLs with insufficient resources, see CSCsq63926 in the Bug Toolkit.

For example, if you apply this ACL to an interface:

```
permit tcp source source-wildcard destination destination-wildcard range 5 60
permit tcp source source-wildcard destination destination-wildcard range 15 160
permit tcp source source-wildcard destination destination-wildcard range 115 1660
permit tcp source source-wildcard destination destination-wildcard
```

And if this message appears:

```
ACLMGR-2-NOVMR: Cannot generate hardware representation of access list [chars]
```

The flag-related operators are not available. To avoid this issue,

- Move the fourth ACE before the first ACE by using **ip access-list resequence** global configuration command:

```
permit tcp source source-wildcard destination destination-wildcard
permit tcp source source-wildcard destination destination-wildcard range 5 60
permit tcp source source-wildcard destination destination-wildcard range 15 160
permit tcp source source-wildcard destination destination-wildcard range 115 1660
```

or

- Rename the ACL with a name or number that alphanumerically precedes the other ACLs (for example, rename ACL 79 to ACL 1).

You can now apply the first ACE in the ACL to the interface. The switch allocates the ACE to available mapping bits in the Opselect index and then allocates flag-related operators to use the same bits in the TCAM.

IPv4 ACL Configuration Examples

This section provides examples of configuring and applying IPv4 ACLs. For detailed information about compiling ACLs, see the *Cisco IOS Security Configuration Guide, Release 12.4* and to the “Configuring IP Services” section in the “IP Addressing and Services” chapter of the *Cisco IOS IP Configuration Guide, Release 12.4*.

This example uses a standard ACL to allow a port access to a specific Internet host with the address 172.20.128.64.

```
Switch(config)# access-list 6 permit 172.20.128.64 0.0.0
Switch(config)# end
Switch# show access-lists
Standard IP access list 6
 10 permit 172.20.128.64 wildcard bits 0.0.0.0
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip access-group 6 in
```

This example uses an extended ACL to deny to a port traffic coming from port 80 (HTTP). It permits all other types of traffic.

```
Switch(config)# access-list 106 deny tcp any any eq 80
Switch(config)# access-list 106 permit ip any any
Switch(config)# end
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip access-group 106 in
```

Numbered ACLs

This ACL accepts addresses on network 36.0.0.0 subnets and denies all packets coming from 56.0.0.0 subnets. The ACL is applied to packets entering a port.

```
Switch(config)# access-list 2 permit 36.0.0.0 0.255.255.255
Switch(config)# access-list 2 deny 56.0.0.0 0.255.255.255
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip access-group 2 in
```

Extended ACLs

In this example, 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. Because the secure system of the network always accepts mail connections on port 25, the incoming services are controlled.

```
Switch(config)# access-list 102 permit tcp any 128.88.0.0 0.0.255.255 eq 23
Switch(config)# access-list 102 permit tcp any 128.88.0.0 0.0.255.255 eq 25
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip access-group 102 in
```

Named ACLs

This example creates an extended ACL named *marketing_group*. The *marketing_group* ACL allows any TCP Telnet traffic to the destination address and wildcard 171.69.0.0 0.0.255.255 and denies any other TCP traffic. It permits any other IP traffic.

```
Switch(config)# ip access-list extended marketing_group
Switch(config-ext-nacl)# permit tcp any 171.69.0.0 0.0.255.255 eq telnet
Switch(config-ext-nacl)# deny tcp any any
Switch(config-ext-nacl)# permit ip any any
Switch(config-ext-nacl)# exit
```

The *marketing_group* ACL is applied to incoming traffic on a port.

```
Switch(config)# interface gigabitethernet0/2
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 gigabitethernet0/2
Switch(config-if)# ip access-group strict in
```

Commented IP ACL Entries

In this example of a numbered ACL, the workstation that belongs to Jones is allowed access, and the workstation that belongs 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
```

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

MAC ACLs are supported only when the switch is running the LAN base image.

For more information about the supported non-IP protocols in the **mac access-list extended** command, see 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.

	Command	Purpose
Step 3	{deny permit} {any host source MAC address source MAC address mask} {any host destination MAC address destination MAC address mask} [type mask lsap lsap mask aarp amber dec-spanning decnet-iv diagnostic dsm etype-6000 etype-8042 lat larc-sca mop-console mop-dump msdos mumps netbios vines-echo vines-ip xns-idp 0-65535] [cos cos]	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: <ul style="list-style-type: none"> <i>type mask</i>—An arbitrary EtherType number of a packet with Ethernet II or SNAP encapsulation in decimal, hexadecimal, or octal with optional mask of <i>don't care</i> bits applied to the EtherType before testing for a match. <i>lsap lsap mask</i>—An LSAP number of a packet with IEEE 802.2 encapsulation in decimal, hexadecimal, 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 larc-sca mop-console mop-dump msdos mumps netbios vines-echo vines-ip xns-idp—A non-IP protocol. cos cos—An IEEE 802.1Q cost of service number from 0 to 7 used to set priority.
Step 4	end	Return to privileged EXEC mode.
Step 5	show access-lists [<i>number 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
    10 deny any any decnet-iv
    20 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-IP traffic coming in that interface. When you apply the MAC ACL, consider these guidelines:

- 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 IP packets, and the MAC access list filters non-IP 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. 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* to a port to filter packets entering the port:

```
Switch(config)# interface gigabitethernet0/2
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.

Displaying IPv4 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.

When you use the **ip access-group** interface configuration command to apply ACLs to a Layer 2 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 1-2](#) to display this information.

Table 1-2 *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 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.



CHAPTER 1

Configuring QoS

This chapter describes how to configure quality of service (QoS) by using automatic QoS (auto-QoS) commands or by using standard QoS commands on the Catalyst 2960 and 2960-S switches. With QoS, you can provide preferential treatment to certain types of 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. Unless otherwise noted, the term *switch* refers to a standalone switch and a switch stack.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.



Note

Catalyst 2960-S switches do not support ingress queueing.

You can configure QoS only on physical ports. You configure the QoS settings, such as classification, queueing, and scheduling, and apply the policy map to a port.



Note

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release.

This chapter consists of these sections:

- [Understanding QoS, page 1-2](#)
- [Configuring Auto-QoS, page 1-21](#)



Note

To use auto-QoS, the switch must be running the LAN Base image.

- [Displaying Auto-QoS Information, page 1-37](#)
- [Configuring Standard QoS, page 1-37](#)
- [Displaying Standard QoS Information, page 1-81](#)

The switch supports some of the modular QoS CLI (MQC) commands. For more information about the MQC commands, see the “Modular Quality of Service Command-Line Interface Overview”:
http://www.cisco.com/en/US/docs/ios/12_2/qos/configuration/guide/qcfmcli2.html

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 can also be carried in the Layer 2 frame. These special bits in the Layer 2 frame or a Layer 3 packet are described here and shown in [Figure 1-1](#):

- 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 class of service (CoS) value in the three least-significant bits. On ports configured as Layer 2 ISL trunks, all traffic is in ISL frames.

Layer 2 IEEE 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 IEEE 802.1Q trunks, all traffic is in IEEE 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:



Note To use DSCP, the switch must be running the LAN Base image.

Layer 3 IP packets can carry either an IP precedence value or a Differentiated Services Code Point (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.




Note IPv6 QoS is not supported in this release.

Figure 1-1 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



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 incoming and outgoing traffic.

Basic QoS Model

To implement QoS, the switch must distinguish packets or flow from one another (classify), 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).

**Note**

Catalyst 2960-S switches do not support ingress queuing.

Figure 1-2 shows the basic QoS model. Actions at the ingress port include classifying traffic, policing, marking, queueing, and scheduling:

- Classifying 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 [“Classification” section on page 1-5](#).
- Policing determines whether a packet is in or out of profile by comparing the rate of the incoming 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 [“Policing and Marking” section on page 1-9](#).
- Marking evaluates the policer and configuration information for the action to be taken when a packet is out of profile and determines 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 [“Policing and Marking” section on page 1-9](#).
- 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 weighted tail-drop (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 1-12](#).
- Scheduling services the queues based on their configured shaped round robin (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 1-14](#).

Actions at the egress port include queueing and scheduling:

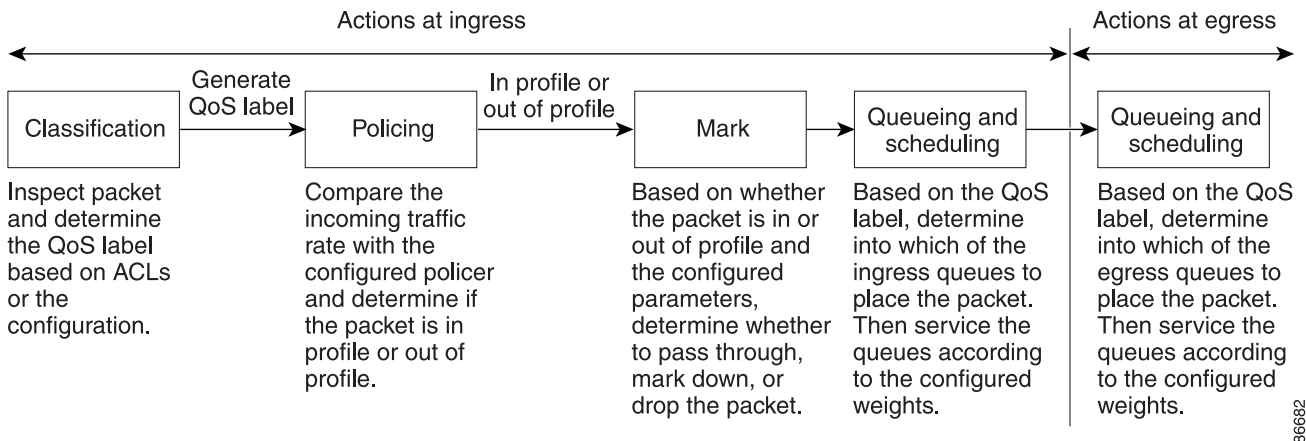
- Queueing evaluates the QoS packet label and the corresponding DSCP or CoS value before selecting which of the four egress queues to use. Because congestion can occur when multiple ingress ports simultaneously send data to an egress port, WTD differentiates traffic classes and subjects the packets to 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 1-12](#).
- Scheduling services the four egress queues based on their configured SRR shared or shaped weights. One of the queues (queue 1) can be the expedited queue, which is serviced until empty before the other queues are serviced.



Note

Catalyst 2960-S switches do not support ingress queueing.

Figure 1-2 Basic QoS Model



Classification

Classification is the process of distinguishing one kind of traffic from another by examining the fields in the packet. Classification is enabled only if QoS is globally enabled on the switch. By default, QoS is globally disabled, so no classification occurs.

During 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 1-3 on page 1-7](#).

You specify which fields in the frame or packet that you want to use to classify incoming traffic. For non-IP traffic, you have these classification options as shown in [Figure 1-3](#):

- Trust the CoS value in the incoming 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 3 least-significant bits of the 1-byte User field. Layer 2 IEEE 802.1Q frame headers carry the CoS value in the 3 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 incoming 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 an internal DSCP value from the CoS-to-DSCP map. The switch uses the internal DSCP value to generate a CoS value representing the priority of the traffic.
- Perform the classification based on a configured Layer 2 MAC access control list (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 CoS value to assign to the incoming frame.

For IP traffic, you have these classification options as shown in [Figure 1-3](#):

- Trust the DSCP value in the incoming packet (configure the port to trust DSCP), and assign the same DSCP value to the packet. The IETF defines the 6 most-significant bits of the 1-byte ToS field as the DSCP. The priority represented by a particular DSCP value is configurable. 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 by using the configurable DSCP-to-DSCP-mutation map.

- Trust the IP precedence value in the incoming packet (configure the port to trust IP precedence), and generate a DSCP value for the packet by using the configurable IP-precedence-to-DSCP map. The IP Version 4 specification defines the 3 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 incoming packet, and generate a DSCP value for the packet by using 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 CoS value to assign to the incoming frame.

For information on the maps described in this section, see the [“Mapping Tables” section on page 1-11](#). For configuration information on port trust states, see the [“Configuring Classification Using Port Trust States” section on page 1-43](#).


After classification, the packet is sent to the policing, marking, and the ingress queueing and scheduling stages.



Note

Catalyst 2960-S switches do not support ingress queueing.

Figure 1-3 Classification Flowchart



86834

Classification Based on QoS ACLs

**Note**

If the switch is running the LAN Lite image, you can configure ACLs, but you cannot attach them to physical interfaces. You can attach them to VLAN interfaces to filter traffic to the CPU.

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 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.

After a traffic class has been defined with the ACL, you can attach a policy to it. A policy might contain multiple classes with actions specified for each one of them. A policy might include commands to classify the class as a particular aggregate (for example, assign a DSCP) or rate-limit the class. This policy is then attached to a particular port on which it becomes effective.

You implement IP ACLs to classify IP traffic by using the **access-list** global configuration command; you implement Layer 2 MAC ACLs to classify non-IP traffic by using the **mac access-list extended** global configuration command. For configuration information, see the [“Configuring a QoS Policy” section on page 1-50](#).

Classification Based on Class Maps and Policy Maps

**Note**

To use policy maps, the switch must be running the LAN Base image.

A class map is a mechanism that you use to name 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. After a packet is matched against the class-map criteria, you further classify it through the use of a policy map.

A policy map specifies which traffic class to act on. Actions can include trusting the 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. Before a policy map can be effective, you must attach it to a port.

You create a class map by using the **class-map** global configuration command or the **class** policy-map configuration command. You should use the **class-map** command when the map is shared among many ports. 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.

You can configure a default class by using the **class class-default** policy-map configuration command (not supported on Catalyst 2960-S). Unclassified traffic (traffic that does not meet the match criteria specified in the traffic classes) is treated as default traffic.

You create and name a policy map by using the **policy-map** global configuration command. When you enter this command, the switch enters the policy-map configuration mode. In this mode, you specify the actions to take on a specific traffic class by using the **class**, **trust**, or **set** policy-map configuration and policy-map class configuration commands.

The policy map can contain the **police** and **police aggregate** policy-map class configuration commands, which define the policer, the bandwidth limitations of the traffic, and the action to take if the limits are exceeded.

To enable the policy map, you attach it to a port by using the **service-policy** interface configuration command.

For more information, see the “Policing and Marking” section on page 1-9. For configuration information, see the “Configuring a QoS Policy” section on page 1-50.

Policing and Marking



Note

To use policing and marking, the switch must be running the LAN Base image.

After a packet is classified and has a DSCP-based or CoS-based QoS label assigned to it, the policing and marking process can begin as shown in [Figure 1-4](#).

Policing involves creating a policer that specifies the bandwidth limits for the traffic. Packets that exceed the limits are *out of profile* or *nonconforming*. Each policer 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 1-11. Marked-down packets use the same queues as the original QoS label to prevent packets in a flow from getting out of order.



Note

All traffic, regardless of whether it is bridged or routed, is subjected to a policer, if one is configured. 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 configure policing on a physical port. For more information about configuring policing on physical ports, see the “[Policing on Physical Ports](#)” section on page 1-10.

After you configure the policy map and policing actions, attach the policy to an ingress port by using the **service-policy** interface configuration command. For configuration information, see the “[Classifying, Policing, and Marking Traffic on Physical Ports by Using Policy Maps](#)” section on page 1-55 and the “[Classifying, Policing, and Marking Traffic by Using Aggregate Policers](#)” section on page 1-60.

**Note**

Catalyst 2960-S switches do not support ingress queueing.

Policing on Physical Ports

In policy maps on physical ports, you can create these types of policers:

- Individual—QoS applies the bandwidth limits specified in the policer separately to each matched traffic class. You configure this type of policer within a policy map by using the **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.

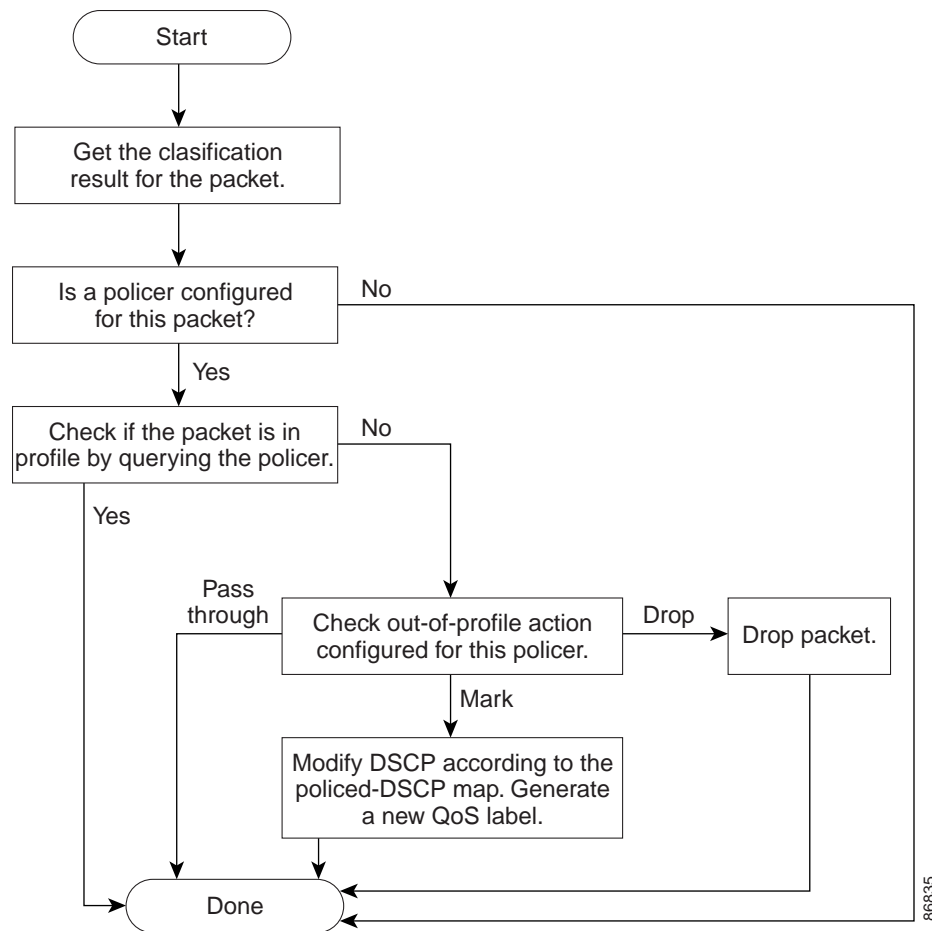
Policing uses a token-bucket algorithm. As each frame is received by the switch, a token is added to the bucket. The bucket has a hole in it and leaks at a rate that you specify as the average traffic rate in bits per second. Each time a token is added to the bucket, the switch verifies that there is 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-b/s), 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 transmitted 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.

Figure 1-4 shows the policing and marking process.

Figure 1-4 Policing and Marking Flowchart on Physical Ports



Mapping Tables



Note To use mapping tables, the switch must be running the LAN Base image.

During QoS 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 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.



Note Catalyst 2960-S switches do not support ingress queueing.

- During 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 packet and selects the queue through the DSCP input and output queue threshold maps or through the CoS input and output queue threshold maps. In addition to an ingress or an egress queue, the QoS label also identifies the WTD threshold value. 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 incoming 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.

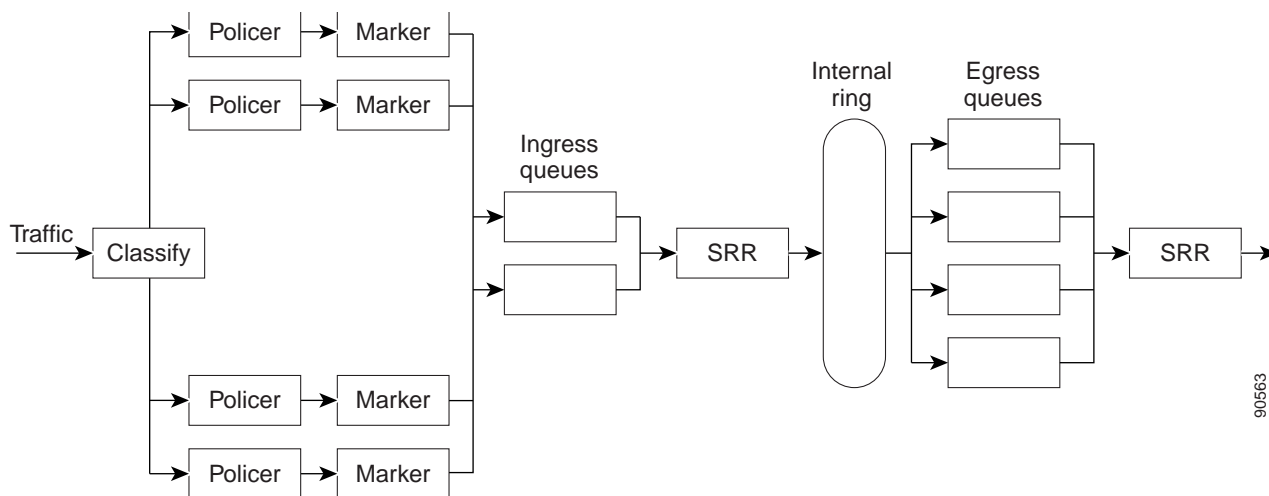
For configuration information, see the [“Configuring DSCP Maps”](#) section on page 1-63.

For information about the DSCP and CoS input queue threshold maps, see the [“Queueing and Scheduling on Ingress Queues”](#) section on page 1-15. For information about the DSCP and CoS output queue threshold maps, see the [“Queueing and Scheduling on Egress Queues”](#) section on page 1-17.

Queueing and Scheduling Overview

The switch has queues at specific points to help prevent congestion as shown in [Figure 1-5](#).

Figure 1-5 Ingress and Egress Queue Location



90563

Because the total inbound 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, outbound queues are located after the internal ring.

Weighted Tail Drop



Note

To use WTD, the switch must be running the LAN Base image.



Note

Catalyst 2960-S switches do not support ingress queueing.

Both the ingress and egress queues use an enhanced version of the tail-drop congestion-avoidance mechanism called weighted tail drop (WTD). WTD is implemented on queues to manage the queue lengths and to provide drop precedences for different traffic classifications.

As a frame is enqueued 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.


Each queue has three threshold values. The QoS label determines which of the three threshold values is subjected to the frame. Of the three thresholds, two are configurable (explicit) and one is not (implicit).

Figure 1-6 shows an example of WTD operating on a queue whose size is 1000 frames. Three drop percentages are configured: 40 percent (400 frames), 60 percent (600 frames), and 100 percent (1000 frames). 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 1-6 *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 1-69, the “[Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set](#)” section on page 1-74, and the “[Mapping DSCP or CoS Values to an Egress Queue and to a Threshold ID](#)” section on page 1-76.

SRR Shaping and Sharing

Both the ingress and egress queues are serviced by SRR, which controls the rate at which packets are sent. On the ingress queues, SRR sends packets to the stack or internal ring. On the egress queues, SRR sends packets to the egress port.

**Note**

Catalyst 2960-S switches do not support ingress queueing.

You can configure SRR on egress queues for sharing or for shaping. However, for ingress queues, sharing is the default mode, and it is the only mode supported.

In shaped mode, the egress 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. Shaping and sharing is configured per interface. Each interface can be uniquely configured.

For more information, see the [“Allocating Bandwidth Between the Ingress Queues”](#) section on page 1-72, the [“Configuring SRR Shaped Weights on Egress Queues”](#) section on page 1-78, and the [“Configuring SRR Shared Weights on Egress Queues”](#) section on page 1-79.

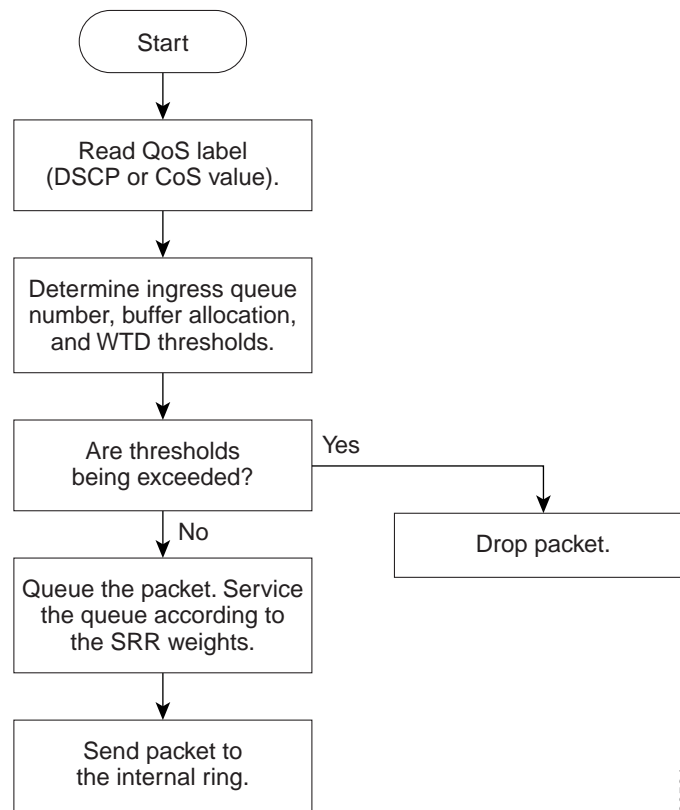
Queueing and Scheduling on Ingress Queues



Note Catalyst 2960-S switches do not support ingress queueing.

Figure 1-7 shows the queueing and scheduling flowchart for ingress ports.

Figure 1-7 Queueing and Scheduling Flowchart for Ingress Ports



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 1-1](#) describes the queues.



Note Catalyst 2960-S switches do not support ingress queueing.

Table 1-1 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 stack 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 and stack 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



Note

Catalyst 2960-S switches do not support ingress queueing.

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-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 1-13.

Buffer and Bandwidth Allocation



Note

Catalyst 2960-S switches do not support ingress queueing.

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.

Priority Queueing



Note

Catalyst 2960-S switches do not support ingress 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 stack or 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 DSCPs or CoSs 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 1-69](#).

Queueing and Scheduling on Egress Queues


[Figure 1-8](#) shows the queueing and scheduling flowchart for egress ports.



Note

If the expedite queue is enabled, SRR services it until it is empty before servicing the other three queues.

Figure 1-8 Queueing and Scheduling Flowchart for Egress Ports




90565

Each port supports four egress queues, one of which (queue 1) can be the egress expedite queue. These queues are configured by a queue-set. All traffic leaving an egress port flows through one of these four queues and is subjected to a threshold based on the QoS label assigned to the packet.

Figure 1-9 shows the egress queue 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 the target queue has not 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 1-9 Egress Queue Buffer Allocation



86695

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 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 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 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. You map a port to queue-set by using the **queue-set *qset-id*** interface configuration command.

Modify the queue-set configuration to change the WTD threshold percentages. For more information about how WTD works, see the [“Weighted Tail Drop” section on page 1-13](#).

Shaped or Shared Mode

SRR services each queue-set in shared or shaped mode. You assign shared or shaped weights to the port by using the **srr-queue bandwidth share** *weight1 weight2 weight3 weight4* or the **srr-queue bandwidth shape** *weight1 weight2 weight3 weight4* interface configuration commands. For an explanation of the differences between shaping and sharing, see the [“SRR Shaping and Sharing” section on page 1-14](#).

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.

All four queues participate in the SRR unless the expedite queue is enabled, in which case the first bandwidth weight is ignored and is not used in the ratio calculation. The expedite queue is a priority queue, and it is serviced until empty before the other queues are serviced. You enable the expedite queue by using the **priority-queue out** interface configuration command.

You can combine the commands described in this section to prioritize traffic by placing packets with particular DSCPs or CoSs 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 Characteristics” section on page 1-74](#).



Note

The egress queue default settings are suitable for most situations. You should change them only when you have a thorough understanding of the egress queues and if these settings do not meet your QoS solution.

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, 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.
- During 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.
- 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 incoming 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 incoming 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.

Configuring Auto-QoS

**Note**

To use auto-QoS, the switch must be running the LAN Base image.

**Note**

Catalyst 2960-S switches do not support ingress queuing.

You can use the auto-QoS feature to simplify the deployment of QoS features. Auto-QoS determines the network design and enables QoS configurations so that the switch can prioritize different traffic flows. It uses the ingress and egress queues instead of using the default (disabled) QoS behavior. The switch offers best-effort service to each packet, regardless of the packet contents or size, and sends it from a single queue.

When you enable auto-QoS, it automatically classifies traffic based on the traffic type and ingress packet label. The switch uses the classification results to choose the appropriate egress queue.

You use auto-QoS commands to identify ports connected to these Cisco devices:

- Cisco IP Phones
- Devices running the Cisco SoftPhone application
- Cisco TelePresence
- Cisco IP Camera

You also use the commands to identify ports that receive trusted traffic through an uplink. Auto-QoS then performs these functions:

- Detects the presence or absence of auto-QoS devices through conditional trusted interfaces.
- Configures QoS classification
- Configures egress queues

These sections contain this configuration information:

- [Generated Auto-QoS Configuration, page 1-22](#)
- [Effects of Auto-QoS on the Configuration, page 1-34](#)
- [Auto-QoS Configuration Guidelines, page 1-34](#)
- [Enabling Auto-QoS, page 1-35](#)

Generated Auto-QoS Configuration

By default, auto-QoS is disabled on all ports. Packets are not modified--the CoS, DSCP and IP precedence values in the packet are not changed.



Note

Catalyst 2960-S switches do not support ingress queuing.

When you enable the auto-QoS feature on the first port of the interface:

- Ingress packet label is used to categorize traffic, to assign packet labels, and to configure the ingress and egress queues.
- QoS is globally enabled (**mls qos** global configuration command), and other global configuration commands are automatically generated. (See [Table 1-5](#).)
- Switch enables the trusted boundary feature and uses the Cisco Discovery Protocol (CDP) to detect the presence of a supported device.
- Policing is used to determine whether a packet is in or out of profile and specifies the action on the packet.

VOIP Device Specifics

- When you enter the **auto qos voip cisco-phone** command on a port at the network edge connected to a Cisco IP Phone, the switch enables the trusted boundary feature. If the packet does not have a DSCP value of 24, 26, or 46 or is out of profile, the switch changes the DSCP value to 0. When there is no Cisco IP Phone, the ingress classification is set to not trust the QoS label in the packet. The policing is applied to the traffic matching the policy-map classification before the switch enables the trust boundary feature.
- When you enter the **auto qos voip cisco-softphone** interface configuration command on a port at the network edge that is connected to a device running the Cisco SoftPhone, the switch uses policing to determine whether a packet is in or out of profile and to specify the action on the packet. If the packet does not have a DSCP value of 24, 26, or 46 or is out of profile, the switch changes the DSCP value to 0.
- When you enter the **auto qos voip trust** interface configuration command on a port connected to the network interior, 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 ingress and egress queues on the port according to the settings in [Table 1-2](#) and [Table 1-3](#).

Table 1-2 Traffic Types, Packet Labels, and Queues

	VoIP ¹ Data Traffic	VoIP Control Traffic	Routing Protocol Traffic	STP BPDU Traffic	Real-Time Video Traffic	All Other Traffic
DSCP	46	24, 26	48	56	34	–
CoS	5	3	6	7	3	–

Table 1-2 Traffic Types, Packet Labels, and Queues (continued)

	VoIP ¹ Data Traffic	VoIP Control Traffic	Routing Protocol Traffic	STP BPDU Traffic	Real-Time Video Traffic	All Other Traffic
CoS-to-Ingress Queue Map	4, 5 (queue 2)					0, 1, 2, 3, 6, 7(queue 1)
CoS-to-Egress Queue Map	4, 5 (queue 1)	2, 3, 6, 7 (queue 2)			0 (queue 3)	2 (queue 3) 0, 1 (queue 4)

1. VoIP = voice over IP

Table 1-3 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, 2, 3, 6, 7	70 percent	90 percent
Priority	2	4, 5	30 percent	10 percent

Table 1-4 Auto-QoS Configuration for the Egress Queues

Egress Queue	Queue Number	CoS-to-Queue Map	Queue Weight (Bandwidth)	Queue (Buffer) Size for Gigabit-Capable Ports	Queue (Buffer) Size for 10/100 Ethernet Ports
Priority	1	4, 5	up to 100 percent	25 percent	15 percent
SRR shared	2	2, 3, 6, 7	10 percent	25 percent	25 percent
SRR shared	3	0	60 percent	25 percent	40 percent
SRR shared	4	1	20 percent	25 percent	20 percent

For information about the trusted boundary feature, see the [“Configuring a Trusted Boundary to Ensure Port Security”](#) section on page 39-42.

When you enable auto-QoS by using the **auto qos voip cisco-phone**, the **auto qos voip cisco-softphone**, or the **auto qos voip trust** interface configuration command, the switch automatically generates a QoS configuration based on the traffic type and ingress packet label and applies the commands listed in [Table 1-5](#) to the port.

Enhanced Auto-QoS for Video, Trust, and Classification



Note

The enhanced auto-QoS feature is not supported on switches running LAN Lite images.

In Cisco IOS Release 12.2(55)SE, auto-QoS is enhanced to support video. Automatic configurations are generated that classify and trust traffic from Cisco TelePresence systems and Cisco IP cameras.

When you configure the **auto qos {video | classify | trust}** enhanced commands on a switch port, this behavior occurs:

- **Auto qos voip** generated commands that you configured on the interface before Cisco IOS Release 12.2(55)SE migrate to the enhanced commands.
- Global values change with the migration of enhanced commands. For a complete list of the generated commands that are applied to the running configuration see [Table 1-5](#).

Auto-QoS Configuration Migration

Auto-QoS configuration migration from legacy auto-QoS to enhanced auto-QoS occurs when:

- A switch is booted with the Cisco IOS Release 12.2(55)SE image and QoS is not enabled. Any video or voice trust configuration on the interface automatically generates enhanced auto-QoS commands.
- A switch is enabled with QoS, these guidelines take effect:
 - If you configure the interface for conditional trust on a voice device, only the legacy auto-QoS VoIP configuration is generated.
 - If you configure the interface for conditional trust on a video device, the enhanced auto-QoS configuration is generated.
 - If you configure the interface with classification or conditional trust based on the new interface auto-QoS commands, enhanced auto-QoS configuration is generated.
- Auto-QoS migration happens after a new device is connected when the **auto qos srnd4** global configuration command is enabled.



Note

If an interface previously configured with legacy auto-QoS migrates to enhanced auto-QoS, voice commands and configuration are updated to match the new global QoS commands.

Auto-QoS configuration migration from enhanced auto-QoS to legacy auto-QoS can occur only when you disable all existing auto-QoS configurations from the interface.



Note

Catalyst 2960-S switches do not support ingress queuing.

Global Auto-QoS Configuration

Table 1-5 Generated Auto-QoS Configuration

Description	Automatically Generated Command {voip}	Enhanced Automatically Generated Command {Video Trust Classify}
The switch automatically enables standard QoS and configures the CoS-to-DSCP map (maps CoS values in incoming packets to a DSCP value).	<pre>Switch(config)# mls qos Switch(config)# mls qos map cos-dscp 0 8 16 26 32 46 48 56</pre>	<pre>Switch(config)# mls qos Switch(config)# mls qos map cos-dscp 0 8 16 24 32 46 48 56</pre>
<p>The switch automatically maps CoS values to an ingress queue and to a threshold ID.</p> <p>Note Catalyst 2960-S switches do not support ingress queueing.</p>	<pre>Switch(config)# no mls qos srr-queue input cos-map Switch(config)# mls qos srr-queue input cos-map queue 1 threshold 2 1 Switch(config)# mls qos srr-queue input cos-map queue 1 threshold 3 0 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</pre>	<pre>Switch(config)# no mls qos srr-queue input cos-map Switch(config)# mls qos srr-queue input cos-map queue 1 threshold 2 3 Switch(config)# mls qos srr-queue input cos-map queue 1 threshold 3 6 7 Switch(config)# mls qos srr-queue input cos-map queue 2 threshold 1 4</pre>
The switch automatically maps CoS values to an egress queue and to a threshold ID.	<pre>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</pre>	<pre>Switch(config)# no mls qos srr-queue output cos-map Switch(config)# mls qos srr-queue output cos-map queue 1 threshold 3 4 5 Switch(config)# mls qos srr-queue output cos-map queue 2 threshold 3 6 7 Switch(config)# mls qos srr-queue output cos-map queue 2 threshold 1 2 Switch(config)# mls qos srr-queue output cos-map queue 2 threshold 2 3 Switch(config)# mls qos srr-queue output cos-map queue 3 threshold 3 0 Switch(config)# mls qos srr-queue output cos-map queue 4 threshold 3 1</pre>

Table 1-5 Generated Auto-QoS Configuration (continued)

Description	Automatically Generated Command {voip}	Enhanced Automatically Generated Command {Video Trust Classify}
<p>The switch automatically maps DSCP values to an ingress queue and to a threshold ID.</p>	<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 33 34 35 36 37 38 39 48 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 2 49 50 51 52 53 54 55 56 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 2 57 58 59 60 61 62 63 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 3 24 25 26 27 28 29 30 31 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 3 40 41 42 43 44 45 46 47</pre>	<pre>Switch(config)# no mls qos srr-queue input dscp-map Switch(config)# mls qos srr-queue input dscp-map queue 1 threshold 2 24 49 50 51 52 53 54 55 56 Switch(config)# mls qos srr-queue input dscp-map queue 1 threshold 3 57 58 59 60 61 62 63 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 3 32 33 40 41 42 43 44 45 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 3 46 47</pre>
<p>Note Catalyst 2960-S switches do not support ingress queueing.</p>		

Table 1-5 Generated Auto-QoS Configuration (continued)

Description	Automatically Generated Command {voip}	Enhanced Automatically Generated Command {Video Trust Classify}
The switch automatically maps DSCP values to an egress queue 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 26 27 28 29 30 31 Switch(config)# mls qos srr-queue output dscp-map queue 2 threshold 3 48 49 50 51 52 53 54 55 Switch(config)# mls qos srr-queue output dscp-map queue 2 threshold 3 56 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 32 33 34 35 36 37 38 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>	<pre>Switch(config)# no mls qos srr-queue output dscp-map Switch(config)# mls qos srr-queue output dscp-map queue 1 threshold 3 32 33 40 41 42 43 44 45 46 47 Switch(config)# mls qos srr-queue output dscp-map queue 2 threshold 1 16 17 18 19 20 21 22 23 Switch(config)# mls qos srr-queue output dscp-map queue 2 threshold 1 26 27 28 29 30 31 34 35 36 37 38 39 Switch(config)# mls qos srr-queue output dscp-map queue 2 threshold 2 24 Switch(config)# mls qos srr-queue output dscp-map queue 2 threshold 3 48 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 0 1 2 3 4 5 6 7 Switch(config)# mls qos srr-queue output dscp-map queue 4 threshold 1 8 9 11 13 15 Switch(config)# mls qos srr-queue output dscp-map queue 4 threshold 2 10 12 14</pre>

Table 1-5 Generated Auto-QoS Configuration (continued)

Description	Automatically Generated Command {voip}	Enhanced Automatically Generated Command{Video Trust Classify}
<p>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.</p> <p>Note Catalyst 2960-S switches do not support ingress queuing.</p>	<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)# mls qos srr-queue input threshold 1 8 16 Switch(config)# mls qos srr-queue input threshold 2 34 66 Switch(config)# mls qos srr-queue input buffers 67 33</pre>	<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 70 30 Switch(config)# mls qos srr-queue input threshold 1 80 90 Switch(config)# mls qos srr-queue input priority-queue 2 bandwidth 30</pre>
<p>The switch automatically configures the egress queue buffer sizes. It configures the bandwidth and the SRR mode (shaped or shared) on the egress queues mapped to the port.</p>	<pre>Switch(config)# mls qos queue-set output 1 threshold 1 138 138 92 138 Switch(config)# mls qos queue-set output 1 threshold 2 138 138 92 400 Switch(config)# mls qos queue-set output 1 threshold 3 36 77 100 318 Switch(config)# mls qos queue-set output 1 threshold 4 20 50 67 400 Switch(config)# mls qos queue-set output 2 threshold 1 149 149 100 149 Switch(config)# mls qos queue-set output 2 threshold 2 118 118 100 235 Switch(config)# mls qos queue-set output 2 threshold 3 41 68 100 272 Switch(config)# mls qos queue-set output 2 threshold 4 42 72 100 242 Switch(config)# mls qos queue-set output 1 buffers 10 10 26 54 Switch(config)# mls qos queue-set output 2 buffers 16 6 17 61 Switch(config-if)# priority-queue out Switch(config-if)# srr-queue bandwidth share 10 10 60 20</pre>	<pre>Switch(config)# mls qos queue-set output 1 threshold 2 100 100 50 200 Switch(config)# mls qos queue-set output 1 threshold 2 125 125 100 400 Switch(config)# mls qos queue-set output 1 threshold 3 100 100 100 400 Switch(config)# mls qos queue-set output 1 threshold 4 60 150 50 200 Switch(config)# mls qos queue-set output 1 buffers 15 25 40 20</pre>

Auto-QoS Generated Configuration For VoIP Devices

Table 1-6 Generated Auto-QoS Configuration

Description	Automatically Generated Command {voip}
The switch automatically enables standard QoS and configures the CoS-to-DSCP map (maps CoS values in incoming packets to a DSCP value).	<pre>Switch(config)# mls qos Switch(config)# mls qos map cos-dscp 0 8 16 26 32 46 48 56</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. Note Catalyst 2960-S switches do not support ingress queueing.	<pre>Switch(config)# no mls qos srr-queue input cos-map Switch(config)# mls qos srr-queue input cos-map queue 1 threshold 2 1 Switch(config)# mls qos srr-queue input cos-map queue 1 threshold 3 0 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</pre>
The switch automatically maps CoS values to an egress queue and to a threshold ID.	<pre>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</pre>
The switch automatically maps DSCP values to an ingress queue and to a threshold ID. Note Catalyst 2960-S switches do not support ingress queueing.	<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 33 34 35 36 37 38 39 48 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 2 49 50 51 52 53 54 55 56 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 2 57 58 59 60 61 62 63 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 3 24 25 26 27 28 29 30 31 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 3 40 41 42 43 44 45 46 47</pre>

Table 1-6 Generated Auto-QoS Configuration (continued)

Description	Automatically Generated Command {voip}
The switch automatically maps DSCP values to an egress queue 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 26 27 28 29 30 31 Switch(config)# mls qos srr-queue output dscp-map queue 2 threshold 3 48 49 50 51 52 53 54 55 Switch(config)# mls qos srr-queue output dscp-map queue 2 threshold 3 56 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 32 33 34 35 36 37 38 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. Note Catalyst 2960-S switches do not support ingress queueing.	<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)# mls qos srr-queue input threshold 1 8 16 Switch(config)# mls qos srr-queue input threshold 2 34 66 Switch(config)# mls qos srr-queue input buffers 67 33</pre>
The switch automatically configures the egress queue buffer sizes. It configures the bandwidth and the SRR mode (shaped or shared) on the egress queues mapped to the port.	<pre>Switch(config)# mls qos queue-set output 1 threshold 1 138 138 92 138 Switch(config)# mls qos queue-set output 1 threshold 2 138 138 92 400 Switch(config)# mls qos queue-set output 1 threshold 3 36 77 100 318 Switch(config)# mls qos queue-set output 1 threshold 4 20 50 67 400 Switch(config)# mls qos queue-set output 2 threshold 1 149 149 100 149 Switch(config)# mls qos queue-set output 2 threshold 2 118 118 100 235 Switch(config)# mls qos queue-set output 2 threshold 3 41 68 100 272 Switch(config)# mls qos queue-set output 2 threshold 4 42 72 100 242 Switch(config)# mls qos queue-set output 1 buffers 10 10 26 54 Switch(config)# mls qos queue-set output 2 buffers 16 6 17 61 Switch(config-if)# priority-que out Switch(config-if)# srr-queue bandwidth share 10 10 60 20</pre>

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
```

If you entered the **auto qos voip cisco-softphone** command, the switch automatically creates class maps and policy maps.

```
Switch(config)# mls qos map policed-dscp 24 26 46 to 0
Switch(config)# class-map match-all AutoQoS-VoIP-RTP-Trust
Switch(config-cmap)# match ip dscp ef
Switch(config)# class-map match-all AutoQoS-VoIP-Control-Trust
Switch(config-cmap)# match ip dscp cs3 af31
Switch(config)# policy-map AutoQoS-Police-SoftPhone
Switch(config-pmap)# class AutoQoS-VoIP-RTP-Trust
Switch(config-pmap-c)# set dscp ef
Switch(config-pmap-c)# police 320000 8000 exceed-action policed-dscp-transmit
Switch(config-pmap)# class AutoQoS-VoIP-Control-Trust
Switch(config-pmap-c)# set dscp cs3
Switch(config-pmap-c)# police 32000 8000 exceed-action policed-dscp-transmit
```

After creating the class maps and policy maps, the switch automatically applies the policy map called *AutoQoS-Police-SoftPhone* to an ingress interface on which auto-QoS with the Cisco SoftPhone feature is enabled.

```
Switch(config-if)# service-policy input AutoQoS-Police-SoftPhone
```

If you entered the **auto qos voip cisco-phone** command, the switch automatically creates class maps and policy maps.

```
Switch(config-if)# mls qos trust device cisco-phone
```

If you entered the **auto qos voip cisco-softphone** command, the switch automatically creates class maps and policy maps.

```
Switch(config)# mls qos map policed-dscp 24 26 46 to 0
Switch(config)# class-map match-all AutoQoS-VoIP-RTP-Trust
Switch(config-cmap)# match ip dscp ef
Switch(config)# class-map match-all AutoQoS-VoIP-Control-Trust
Switch(config-cmap)# match ip dscp cs3 af31
Switch(config)# policy-map AutoQoS-Police-CiscoPhone
Switch(config-pmap)# class AutoQoS-VoIP-RTP-Trust
Switch(config-pmap-c)# set dscp ef
Switch(config-pmap-c)# police 320000 8000 exceed-action policed-dscp-transmit
Switch(config-pmap)# class AutoQoS-VoIP-Control-Trust
Switch(config-pmap-c)# set dscp cs3
Switch(config-pmap-c)# police 32000 8000 exceed-action policed-dscp-transmit
```

After creating the class maps and policy maps, the switch automatically applies the policy map called *AutoQoS-Police-SoftPhone* to an ingress interface on which auto-QoS with the Cisco SoftPhone feature is enabled.

```
Switch(config-if)# service-policy input AutoQoS-Police-SoftPhone
```

Auto-QoS Generated Configuration For Enhanced Video, Trust, and Classify Devices

If you entered these enhanced auto-QoS commands, the switch automatically configures the CoS-to-DSCP map (maps CoS values in incoming packets to a DSCP value).

- **auto qos video cts**
- **auto qos video ip-camera**
- **auto qos trust**
- **auto qos trust cos**
- **auto qos trust dscp**

```
Switch(config)# mls qos map cos-dscp 0 8 16 24 32 46 48 56
```

**Note**

No class maps and policy maps are configured.

If you entered the **auto qos classify** command, the switch automatically creates class maps and policy maps.

```
Switch(config)# mls qos map policed-dscp 0 10 18 24 26 46 to 8
Switch(config)# mls qos map cos-dscp 0 8 16 24 32 46 48 56
Switch(config)# class-map match-all AUTOQOS_MULTIENTHANCED_CONF_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-MULTIENTHANCED-CONF
Switch(config)# class-map match-all AUTOQOS_DEFAULT_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-DEFAULT
Switch(config)# class-map match-all AUTOQOS_TRANSACTION_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-TRANSACTIONAL-DATA
Switch(config)# class-map match-all AUTOQOS_SIGNALING_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-SIGNALING
Switch(config)# class-map match-all AUTOQOS_BULK_DATA_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-BULK-DATA
Switch(config)# class-map match-all AUTOQOS_SCAVANGER_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-SCAVANGER
Switch(config)# policy-map AUTOQOS-SRND4-CLASSIFY-POLICY
Switch(config-pmap)# class AUTOQOS_MULTIENTHANCED_CONF_CLASS
Switch(config-pmap-c)# set dscp af41
Switch(config-pmap)# class AUTOQOS_BULK_DATA_CLASS
Switch(config-pmap-c)# set dscp af11
Switch(config-pmap)# class AUTOQOS_TRANSACTION_CLASS
Switch(config-pmap-c)# set dscp af21
Switch(config-pmap)# class AUTOQOS_SCAVANGER_CLASS
Switch(config-pmap-c)# set dscp cs1
Switch(config-pmap)# class AUTOQOS_SIGNALING_CLASS
Switch(config-pmap-c)# set dscp cs3
Switch(config-pmap)# class AUTOQOS_DEFAULT_CLASS
Switch(config-pmap-c)# set dscp default
;
Switch(config-if)# service-policy input AUTOQOS-SRND4-CLASSIFY-POLICY
```

If you entered the **auto qos classify police** command, the switch automatically creates class maps and policy maps.

```
Switch(config)# mls qos map policed-dscp 0 10 18 to 8
Switch(config)# mls qos map cos-dscp 0 8 16 24 32 46 48 56
Switch(config)# class-map match-all AUTOQOS_MULTIENTHANCED_CONF_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-MULTIENTHANCED-CONF
Switch(config)# class-map match-all AUTOQOS_DEFAULT_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-DEFAULT
Switch(config)# class-map match-all AUTOQOS_TRANSACTION_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-TRANSACTIONAL-DATA
Switch(config)# class-map match-all AUTOQOS_SIGNALING_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-SIGNALING
Switch(config)# class-map match-all AUTOQOS_BULK_DATA_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-BULK-DATA
Switch(config)# class-map match-all AUTOQOS_SCAVANGER_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-SCAVANGER
Switch(config)# policy-map AUTOQOS-SRND4-CLASSIFY-POLICE-POLICY
Switch(config-pmap)# class AUTOQOS_MULTIENTHANCED_CONF_CLASS
Switch(config-pmap-c)# set dscp af41
Switch(config-pmap-c)# police 5000000 8000 exceed-action drop
Switch(config-pmap)# class AUTOQOS_BULK_DATA_CLASS
Switch(config-pmap-c)# set dscp af11
Switch(config-pmap-c)# police 10000000 8000 exceed-action policed-dscp-transmit
Switch(config-pmap)# class AUTOQOS_TRANSACTION_CLASS
```

```

Switch(config-pmap-c)# set dscp af21
Switch(config-pmap-c)# police 10000000 8000 exceed-action policed-dscp-transmit
Switch(config-pmap)# class AUTOQOS_SCAVANGER_CLASS
Switch(config-pmap-c)# set dscp cs1
Switch(config-pmap-c)# police 10000000 8000 exceed-action drop
Switch(config-pmap)# class AUTOQOS_SIGNALING_CLASS
Switch(config-pmap-c)# set dscp cs3
Switch(config-pmap-c)# police 32000 8000 exceed-action drop
Switch(config-pmap)# class AUTOQOS_DEFAULT_CLASS
Switch(config-pmap-c)# set dscp default
Switch(config-pmap-c)# police 10000000 8000 exceed-action policed-dscp-transmit
;
Switch(config-if)# service-policy input AUTOQOS-SRND4-CLASSIFY-POLICE-POLICY

```

This is the enhanced configuration for the **auto qos voip cisco-phone** command:

```

Switch(config)# mls qos map policed-dscp 0 10 18 to 8
Switch(config)# mls qos map cos-dscp 0 8 16 24 32 46 48 56
Switch(config)# class-map match-all AUTOQOS_VOIP_DATA_CLASS
Switch(config-cmap)# match ip dscp ef
Switch(config)# class-map match-all AUTOQOS_DEFAULT_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-DEFAULT
Switch(config)# class-map match-all AUTOQOS_VOIP_SIGNAL_CLASS
Switch(config-cmap)# match ip dscp cs3
Switch(config)# policy-map AUTOQOS-SRND4-CISCOPHONE-POLICY
Switch(config-pmap)# class AUTOQOS_VOIP_DATA_CLASS
Switch(config-pmap-c)# set dscp ef
Switch(config-pmap-c)# police 128000 8000 exceed-action policed-dscp-transmit
Switch(config-pmap)# class AUTOQOS_VOIP_SIGNAL_CLASS
Switch(config-pmap-c)# set dscp cs3
Switch(config-pmap-c)# police 32000 8000 exceed-action policed-dscp-transmit
Switch(config-pmap)# class AUTOQOS_DEFAULT_CLASS
Switch(config-pmap-c)# set dscp default
Switch(config-pmap-c)# police 10000000 8000 exceed-action policed-dscp-transmit
;
Switch(config-if)# service-policy input AUTOQOS-SRND4-CISCOPHONE-POLICY

```

This is the enhanced configuration for the **auto qos voip cisco-softphone** command:

```

Switch(config)# mls qos map policed-dscp 0 10 18 to 8
Switch(config)# mls qos map cos-dscp 0 8 16 24 32 46 48 56
Switch(config)# class-map match-all AUTOQOS_MULTIHANCED_CONF_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-MULTIHANCED-CONF
Switch(config)# class-map match-all AUTOQOS_VOIP_DATA_CLASS
Switch(config-cmap)# match ip dscp ef
Switch(config)# class-map match-all AUTOQOS_DEFAULT_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-DEFAULT
Switch(config)# class-map match-all AUTOQOS_TRANSACTION_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-TRANSACTIONAL-DATA
Switch(config)# class-map match-all AUTOQOS_VOIP_SIGNAL_CLASS
Switch(config-cmap)# match ip dscp cs3
Switch(config)# class-map match-all AUTOQOS_SIGNALING_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-SIGNALING
Switch(config)# class-map match-all AUTOQOS_BULK_DATA_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-BULK-DATA
Switch(config)# class-map match-all AUTOQOS_SCAVANGER_CLASS
Switch(config-cmap)# match access-group name AUTOQOS-ACL-SCAVANGER

Switch(config)# policy-map AUTOQOS-SRND4-SOFTPHONE-POLICY
Switch(config-pmap)# class AUTOQOS_VOIP_DATA_CLASS
Switch(config-pmap-c)# set dscp ef
Switch(config-pmap-c)# police 128000 8000 exceed-action policed-dscp-transmit
Switch(config-pmap)# class AUTOQOS_VOIP_SIGNAL_CLASS

```

```

Switch(config-pmap-c)# set dscp cs3
Switch(config-pmap-c)# police 32000 8000 exceed-action policed-dscp-transmit
Switch(config-pmap)# class AUTOQOS_MULTTIENHANCED_CONF_CLASS
Switch(config-pmap-c)# set dscp af41
Switch(config-pmap-c)# police 5000000 8000 exceed-action drop
Switch(config-pmap)# class AUTOQOS_BULK_DATA_CLASS
Switch(config-pmap-c)# set dscp af11
Switch(config-pmap-c)# police 10000000 8000 exceed-action policed-dscp-transmit
Switch(config-pmap)# class AUTOQOS_TRANSACTION_CLASS
Switch(config-pmap-c)# set dscp af21
Switch(config-pmap-c)# police 10000000 8000 exceed-action policed-dscp-transmit
Switch(config-pmap)# class AUTOQOS_SCAVANGER_CLASS
Switch(config-pmap-c)# set dscp cs1
Switch(config-pmap-c)# police 10000000 8000 exceed-action drop
Switch(config-pmap)# class AUTOQOS_SIGNALING_CLASS
Switch(config-pmap-c)# set dscp cs3
Switch(config-pmap-c)# police 32000 8000 exceed-action drop
Switch(config-pmap)# class AUTOQOS_DEFAULT_CLASS
Switch(config-pmap-c)# set dscp default
;
Switch(config-if)# service-policy input AUTOQOS-SRND4-SOFTPHONE-POLICY

```

Effects of Auto-QoS on the Configuration

When auto-QoS is enabled, the **auto qos** interface configuration commands and the generated global configuration are added 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 to 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:

- Auto-QoS configures the switch for VoIP with Cisco IP Phones on nonrouted and routed ports. Auto-QoS also configures the switch for VoIP with devices running the Cisco SoftPhone application.
- When a device running Cisco SoftPhone is connected to a nonrouted or routed port, the switch supports only one Cisco SoftPhone application per port.
- Beginning with Cisco IOS Release 12.2(40)SE, Auto-QoS VoIP uses the **priority-queue** interface configuration command for an egress interface. You can also configure a policy-map and trust device on the same interface for Cisco IP phones.
- If the switch port was configured by using the **auto qos voip cisco-phone** interface configuration command in Cisco IOS Release 12.2(37)SE or earlier, the auto-QoS generated commands new to Cisco IOS Release 12.2(40)SE are not applied to the port. To have these commands automatically applied, you must remove and then reapply the configuration to the port.

- 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, page 1-34](#).
- After auto-QoS is enabled, do not modify a policy map or aggregate policer that includes *AutoQoS* in its name. If you need to modify the policy map or aggregate policer, make a copy of it, and change the copied policy map or policer. To use this new policy map instead of the generated one, remove the generated policy map from the interface, and apply the new policy map to the interface.
- You can enable auto-QoS on static, dynamic-access, voice VLAN access, and trunk ports.



Note

VLAN-based QoS is not supported on Catalyst 2960-S switches.

- By default, the CDP is enabled on all ports. For auto-QoS to function properly, do not disable the CDP.
- When enabling auto-QoS with a Cisco IP Phone on a routed port, you must assign a static IP address to the IP phone.
- This release supports only Cisco IP SoftPhone Version 1.3(3) or later.
- Connected devices must use Cisco Call Manager Version 4 or later.

Auto-QoS Enhanced Considerations

- The **auto qos srnd4** global configuration command is generated as a result of enhanced auto-QoS configuration.
- If the legacy **auto qos voip** commands are executed on the switch and the **mls qos** command is disabled, the enhanced auto-QoS configuration is generated. Otherwise, legacy auto-QoS commands are executed.

Enabling Auto-QoS

For optimum QoS performance, enable auto-QoS on all the devices in your network.

Beginning in privileged EXEC mode, follow these steps to enable auto-QoS devices 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 video device or the uplink port that is connected to another trusted switch or router in the network interior, and enter interface configuration mode.

	Command	Purpose
Step 3	auto qos voip { cisco-phone cisco-softphone trust } or	Enable auto-QoS. The keywords have these meanings: <ul style="list-style-type: none"> • cisco-phone—If the port is connected to a Cisco IP Phone, the QoS labels of incoming packets are trusted only when the telephone is detected. • cisco-softphone—The port is connected to device running the Cisco SoftPhone feature. • trust—The uplink port is connected to a trusted switch or router, and the VoIP traffic classification
	auto qos video { cts ip-camera } or	Enable auto-QoS for a video device. <ul style="list-style-type: none"> • cts— A port connected to a Cisco Telepresence system. • ip-camera—A port connected to an IP camera. QoS labels of incoming packets are trusted only when the system is detected.
	auto qos classify [police] or	Enable auto-QoS for classification. <ul style="list-style-type: none"> • police—Policing is set up by defining the QoS policy maps and applying them to ports (port-based QoS).
	auto qos trust { cos dscp }	Enable auto-QoS for trusted interfaces. <ul style="list-style-type: none"> • cos—Class of service. • dscp—Differentiated Services Code Point.
Step 4	exit	Return to global configuration mode.
Step 5	interface <i>interface-id</i>	Specify the switch port identified as connected to a trusted switch or router, and enter interface configuration mode.
Step 6	auto qos trust	Enable auto-QoS on the port, and specify that the port is connected to a trusted router or switch.
Step 7	end	Return to privileged EXEC mode.
Step 8	show auto qos interface <i>interface-id</i>	Verify your entries. This command displays the auto-QoS command on the interface on which auto-QoS was enabled. You can use the show running-config privileged EXEC command to display the auto-QoS configuration and the user modifications.

Troubleshooting Auto QoS Commands

To display the QoS commands that are automatically generated when auto-QoS is enabled or disabled, enter the **debug auto qos** privileged EXEC command *before* you enable auto-QoS. For more information, see the **debug autoqos** command in the command reference for this release.

To disable auto-QoS on a port, use the **no** form of the auto qos command interface configuration command, such as **no auto qos voip**. 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).

Displaying Auto-QoS Information

To display the initial auto-QoS configuration, use the **show auto qos [interface *interface-id*]** privileged EXEC command. To display any user changes to that configuration, use the **show running-config** privileged EXEC command. You can compare the **show auto qos** and the **show running-config** command output to identify the user-defined QoS settings.

To display information about the QoS configuration that might be affected by auto-QoS, use one of these commands:

- **show mls qos**
- **show mls qos 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**



Note Catalyst 2960-S switches do not support ingress queueing.

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 contain this configuration information:

- [Default Standard QoS Configuration, page 1-38](#)
- [Standard QoS Configuration Guidelines, page 1-40](#)
- [Enabling QoS Globally, page 1-43](#) (required)
- [Configuring Classification Using Port Trust States, page 1-43](#) (required)
- [Configuring a QoS Policy, page 1-50](#) (required)
- [Configuring DSCP Maps, page 1-63](#) (optional, unless you need to use the DSCP-to-DSCP-mutation map or the policed-DSCP map)

- [Configuring Ingress Queue Characteristics, page 1-69](#) (optional)
- [Configuring Egress Queue Characteristics, page 1-74](#) (optional)

Default Standard QoS Configuration



Note

Catalyst 2960-S switches do not support ingress queuing.

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 and egress queue settings are described in the “[Default Ingress Queue Configuration](#)” section on [page 1-38](#) and the “[Default Egress Queue Configuration](#)” section on [page 1-39](#).

Default Ingress Queue Configuration



Note

Catalyst 2960-S switches do not support ingress queuing.

[Table 1-7](#) shows the default ingress queue configuration when QoS is enabled.

Table 1-7 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 1-8](#) shows the default CoS input queue threshold map when QoS is enabled.

Table 1-8 Default CoS Input Queue Threshold Map

CoS Value	Queue ID–Threshold ID
0–4	1–1

Table 1-8 *Default CoS Input Queue Threshold Map (continued)*

CoS Value	Queue ID–Threshold ID
5	2–1
6, 7	1–1

Table 1-9 shows the default DSCP input queue threshold map when QoS is enabled.

Table 1-9 *Default DSCP Input Queue Threshold Map*

DSCP Value	Queue ID–Threshold ID
0–39	1–1
40–47	2–1
48–63	1–1

Default Egress Queue Configuration

Table 1-10 shows the default egress queue configuration for each queue-set when QoS is enabled. All ports are mapped to queue-set 1. The port bandwidth limit is set to 100 percent and rate unlimited.

Table 1-10 *Default Egress Queue 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	200 percent	100 percent	100 percent
WTD drop threshold 2	100 percent	200 percent	100 percent	100 percent
Reserved threshold	50 percent	50 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 1-11 shows the default CoS output queue threshold map when QoS is enabled.

Table 1-11 *Default CoS Output Queue Threshold Map*

CoS Value	Queue ID–Threshold ID
0, 1	2–1
2, 3	3–1
4	4–1

Table 1-11 *Default CoS Output Queue Threshold Map*

CoS Value	Queue ID–Threshold ID
5	1–1
6, 7	4–1

[Table 1-12](#) shows the default DSCP output queue threshold map when QoS is enabled.

Table 1-12 *Default DSCP Output Queue Threshold Map*

DSCP Value	Queue ID–Threshold ID
0–15	2–1
16–31	3–1
32–39	4–1
40–47	1–1
48–63	4–1

Default Mapping Table Configuration

The default CoS-to-DSCP map is shown in [Table 1-13 on page 1-63](#).

The default IP-precedence-to-DSCP map is shown in [Table 1-14 on page 1-64](#).

The default DSCP-to-CoS map is shown in [Table 1-15 on page 1-66](#).

The default DSCP-to-DSCP-mutation map is a null map, which maps an incoming DSCP value to the same DSCP value.

The default policed-DSCP map is a null map, which maps an incoming 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 in these sections:

- “[QoS ACL Guidelines](#)” section on page 1-40
- “[Policing Guidelines](#)” section on page 1-42
- “[General QoS Guidelines](#)” section on page 1-42

QoS ACL Guidelines

- If you use QoS ACLs for classification, you can use the **sdm prefer qos** global configuration command to set the Switch Database Management (SDM) feature to the QoS template. SDM configures system resources to support the maximum number of access control entries (ACEs). For more information on the SDM templates, see [Chapter 1, “Configuring SDM Templates.”](#)
- 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 per class map and only one **match** class-map configuration command per class map are supported. The ACL can have multiple ACEs, which match fields against the contents of the packet.
- A trust statement in a policy map requires multiple TCAM entries per ACL line. If an input service policy map contains a trust statement in an ACL, the access-list might be too large to fit into the available QoS TCAM and an error can occur when you apply the policy map to a port. Whenever possible, you should minimize the number of lines in a QoS ACL.

Policing Guidelines



Note To use policing, the switch must be running the LAN Base image.

- The port ASIC device, which controls more than one physical port, supports 256 policers (255 user-configurable policers plus 1 policer reserved for system internal use). The maximum number of user-configurable policers supported per port is 63. 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.
- Only one policer is applied to a packet on an ingress port. Only the average rate and committed burst parameters are configurable.



Note Catalyst 2960-S switches do not support ingress queueing.

- 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.
- If you have EtherChannel ports configured on your switch, you must configure QoS classification, policing, mapping, and queueing on the individual physical ports that comprise the EtherChannel. You must decide whether the QoS configuration should match on all ports in the EtherChannel.
- If you need to modify a policy map of an existing QoS policy, first remove the policy map from all interfaces, and then modify or copy the policy map. After you finish the modification, apply the modified policy map to the interfaces. If you do not first remove the policy map from all interfaces, high CPU usage can occur, which, in turn, can cause the console to pause for a very long time.

General QoS Guidelines

These are general QoS guidelines:

- You configure QoS only on physical ports; there is no support for it at the VLAN level.
- 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.



Note Catalyst 2960-S switches do not support ingress queueing.

- You are likely to lose data when you change queue settings; therefore, try to make changes when traffic is at a minimum.

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 with the default settings described in the “Default Standard QoS Configuration” section on page 1-38, the “Queueing and Scheduling on Ingress Queues” section on page 1-15, and the “Queueing and Scheduling on Egress Queues” section on page 1-17.
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 Classification Using Port Trust States

These sections describe how to classify incoming traffic by using 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 a QoS Policy” section on page 1-50:

- [Configuring the Trust State on Ports within the QoS Domain, page 1-43](#)
- [Configuring the CoS Value for an Interface, page 1-45](#)
- [Configuring a Trusted Boundary to Ensure Port Security, page 1-46](#)
- [Enabling DSCP Transparency Mode, page 1-47](#)
- [Configuring the DSCP Trust State on a Port Bordering Another QoS Domain, page 1-48](#)

Configuring the Trust State on Ports within the QoS Domain

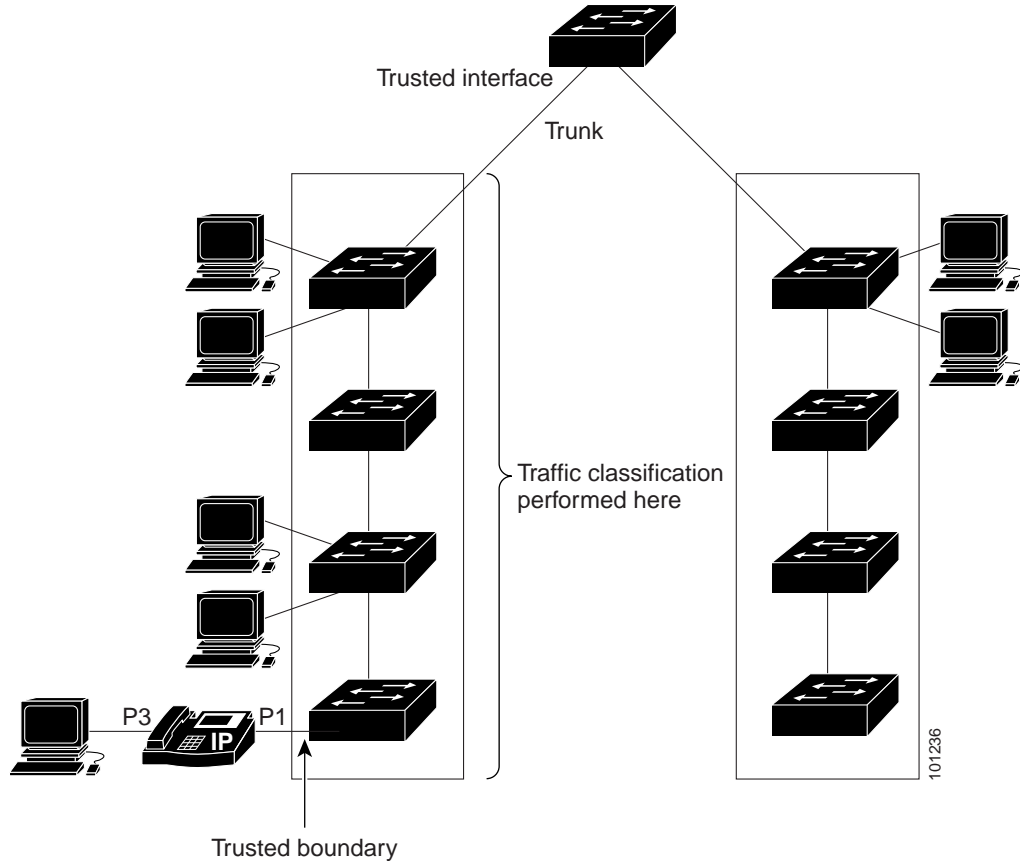


Note

Catalyst 2960-S switches do not support ingress queueing.

Packets entering a QoS domain are classified at the edge of the QoS domain. When the packets are classified at the edge, the switch port within the QoS domain can be configured to one of the trusted states because there is no need to classify the packets at every switch within the QoS domain. [Figure 1-10](#) shows a sample network topology.

Figure 1-10 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 <i>interface-id</i>	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]	<p>Configure the port trust state.</p> <p>By default, the port is not trusted. If no keyword is specified, the default is dscp.</p> <p>The keywords have these meanings:</p> <ul style="list-style-type: none"> • 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. • dscp—Classifies an ingress packet by using the packet DSCP 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 by using the CoS-to-DSCP map. • 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 by using the CoS-to-DSCP map.
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 1-45. For information on how to configure the CoS-to-DSCP map, see the [“Configuring the CoS-to-DSCP Map”](#) section on page 1-63.

Configuring the CoS Value for an Interface



Note Catalyst 2960-S switches do not support ingress queuing.

QoS assigns the CoS value specified with the **mls qos cos** interface configuration command to untagged frames received on trusted and untrusted ports.

Beginning in privileged EXEC mode, follow these steps to define the default CoS value of a port or to assign the default CoS to all incoming packets on the port:

	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.

	Command	Purpose
Step 3	mls qos cos { <i>default-cos</i> override }	Configure the default CoS value for the port. <ul style="list-style-type: none"> 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. Use the override keyword to override the previously configured trust state of the incoming packet and to apply the default port CoS value to the port on all incoming packets. By default, CoS override is disabled. Use the override keyword when all incoming 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 incoming CoS values are assigned the default CoS value configured with this command. If an incoming packet is tagged, the CoS value of the packet is modified with the default CoS of the port at the ingress 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



Note

The trusted boundary feature is supported only on Catalyst 2960-S switches running the LAN base image.

In a typical network, you connect a Cisco IP Phone to a switch port, as shown in [Figure 1-10 on page 1-44](#), 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 IEEE 802.1Q header. The header contains the VLAN information and the class of service (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. Use the **mls qos trust dscp** interface configuration command to configure a routed port to which the telephone is connected to trust the DSCP 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 Cisco 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 Cisco 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 mls qos trust dscp	Configure the switch port to trust the CoS value in traffic received from the Cisco IP Phone. or Configure the routed port to trust the DSCP 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.

Enabling DSCP Transparency Mode

The switch supports the DSCP transparency feature. It affects only the DSCP field of a packet at egress. By default, DSCP transparency is disabled. The switch modifies the DSCP field in an incoming packet, and the DSCP field in the outgoing packet is based on the quality of service (QoS) configuration, including the port trust setting, policing and marking, and the DSCP-to-DSCP mutation map.

If DSCP transparency is enabled by using the **no mls qos rewrite ip dscp** command, the switch does not modify the DSCP field in the incoming packet, and the DSCP field in the outgoing packet is the same as that in the incoming packet.

Regardless of the DSCP transparency configuration, the switch modifies the internal DSCP value of the packet, which the switch uses to generate a class of service (CoS) value that represents the priority of the traffic. The switch also uses the internal DSCP value to select an egress queue and threshold.

Beginning in privileged EXEC mode, follow these steps to enable DSCP transparency on a switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos	Enable QoS globally.
Step 3	no mls qos rewrite ip dscp	Enable DSCP transparency. The switch is configured to not modify the DSCP field of the IP packet.
Step 4	end	Return to privileged EXEC mode.
Step 5	show mls qos interface [<i>interface-id</i>]	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To configure the switch to modify the DSCP value based on the trust setting or on an ACL by disabling DSCP transparency, use the **mls qos rewrite ip dscp** global configuration command.

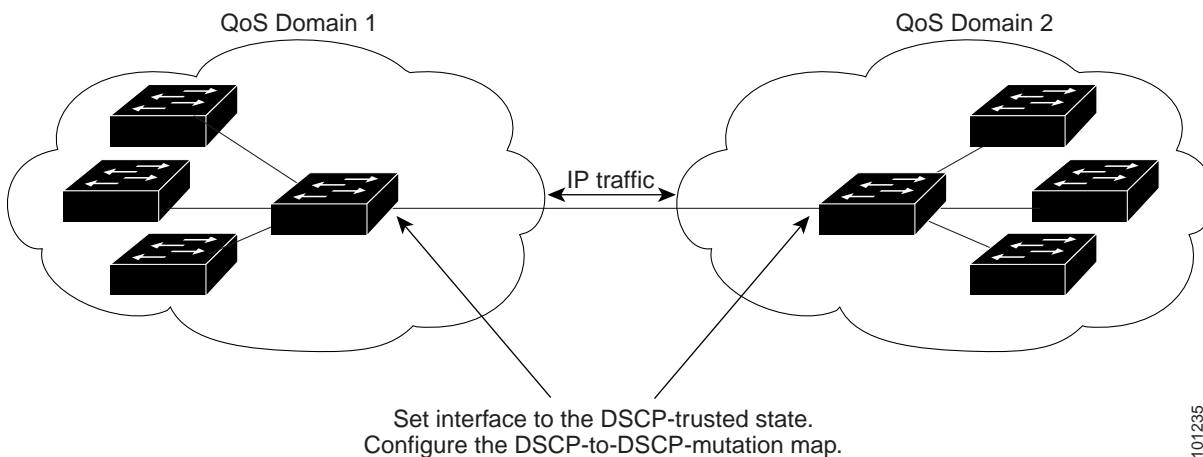
If you disable QoS by using the **no mls qos** global configuration command, the CoS and DSCP values are not changed (the default QoS setting).

If you enter the **no mls qos rewrite ip dscp** global configuration command to enable DSCP transparency and then enter the **mls qos trust [cos | dscp]** interface configuration command, DSCP transparency is still enabled.

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 1-11](#). 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 1-11 DSCP-Trusted State on a Port Bordering Another QoS Domain



101235



Note Catalyst 2960-S switches do not support ingress queueing.

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 incoming DSCP value to the same DSCP value. <ul style="list-style-type: none"> 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 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.
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 *gi0/2-mutation*) so that incoming DSCP values 10 to 13 are mapped to DSCP 30:

```
Switch(config)# mls qos map dscp-mutation gi0/2-mutation 10 11 12 13 to 30
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# mls qos trust dscp
Switch(config-if)# mls qos dscp-mutation gi0/2-mutation
Switch(config-if)# end
```

Configuring a QoS Policy

Configuring a QoS policy typically requires classifying traffic into classes, configuring policies applied to those traffic classes, and attaching policies to ports.



Note

To use policing and marking, the switch must be running the LAN Base image.

For background information, see the [“Classification” section on page 1-5](#) and the [“Policing and Marking” section on page 1-9](#). For configuration guidelines, see the [“Standard QoS Configuration Guidelines” section on page 1-40](#).

These sections describe how to classify, police, and mark traffic. Depending on your network configuration, you must perform one or more of these tasks:

- [Classifying Traffic by Using ACLs, page 1-50](#)
- [Classifying Traffic by Using Class Maps, page 1-53](#)
- [Classifying, Policing, and Marking Traffic on Physical Ports by Using Policy Maps, page 1-55](#)
- [Classifying, Policing, and Marking Traffic by Using Aggregate Policers, page 1-60](#)

Classifying Traffic by Using ACLs

You can classify IP traffic by using IP standard or IP extended ACLs; you can classify non-IP traffic by using Layer 2 MAC ACLs.

Beginning in privileged EXEC mode, follow these steps to create an IP standard ACL for 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>]	<p>Create an IP standard ACL, repeating the command as many times as necessary.</p> <ul style="list-style-type: none"> • 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. <p>Note When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.</p>
Step 3	end	Return to privileged EXEC mode.

	Command	Purpose
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 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>	<p>Create an IP extended ACL, repeating the command as many times as necessary.</p> <ul style="list-style-type: none"> 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>. <p>Note When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.</p>
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 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. <ul style="list-style-type: none"> 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> ffff.ffff.ffff, 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> ffff.ffff.ffff, 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. <p>Note When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.</p>
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 Traffic by Using Class Maps

You use the **class-map** global configuration command to name and to isolate a specific traffic flow (or class) from all other traffic. The class map defines the criteria to use to match against a specific traffic flow to further classify it. Match statements can include criteria such as an ACL, IP precedence values, or DSCP values. The match criterion is defined with one match statement entered within the class-map configuration mode.



Note

You can also create class-maps during policy map creation by using the **class** policy-map configuration command. For more information, see the [“Classifying, Policing, and Marking Traffic on Physical Ports by Using Policy Maps”](#) section on page 1-55.

Beginning in privileged EXEC mode, follow these steps to create a class map and to define the match criterion to classify 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>] or access-list <i>access-list-number</i> {deny permit} <i>protocol source</i> [<i>source-wildcard</i>] <i>destination</i> [<i>destination-wildcard</i>] or mac access-list extended <i>name</i> {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>]	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 Traffic by Using ACLs” section on page 1-50. 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>	<p>Create a class map, and enter class-map configuration mode.</p> <p>By default, no class maps are defined.</p> <ul style="list-style-type: none"> • (Optional) Use the match-all keyword to perform a logical-AND of all matching statements under this class map. All match 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 match criteria must be matched. • For <i>class-map-name</i>, specify the name of the class map. <p>If neither the match-all or match-any keyword is specified, the default is match-all.</p> <p>Note Because only one match command per class map is supported, the match-all and match-any keywords function the same.</p>
Step 4	match { access-group <i>acl-index-or-name</i> ip dscp <i>dscp-list</i> ip precedence <i>ip-precedence-list</i> }	<p>Define the match criterion to classify traffic.</p> <p>By default, no match criterion is defined.</p> <p>Only one match criterion per class map is supported, and only one ACL per class map is supported.</p> <ul style="list-style-type: none"> • 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 incoming 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 incoming 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 policy map, use the **no policy-map** *policy-map-name* global configuration command. 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 ip any any dscp 10
Switch(config)# class-map class1
Switch(config-cmap)# match access-group 103
Switch(config-cmap)# end
Switch#
```

This example shows how to create a class map called *class2*, which matches incoming 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)# end
Switch#
```

This example shows how to create a class map called *class3*, which matches incoming 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)# end
Switch#
```

Classifying, Policing, and Marking Traffic on Physical Ports by Using Policy Maps



Note

To use policing and marking, the switch must be running the LAN Base image.

You can configure a policy map on a physical port that specifies which traffic class to act on. Actions can include trusting the CoS, DSCP, or IP precedence values in the traffic class; setting a specific DSCP or IP precedence value in the traffic class; and specifying the traffic bandwidth limitations for each matched traffic class (policer) and the action to take when the traffic is out of profile (marking).

A policy map also has these characteristics:

- A policy map can contain multiple class statements, each with different match criteria and policers.
- A policy map can contain a predefined default traffic class explicitly placed at the end of the map.
- 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, and whichever is configured last takes affect.

Follow these guidelines when configuring policy maps on physical ports:

- You can attach only one policy map per ingress port.
- If you configure the IP-precedence-to-DSCP map by using the **mls qos map ip-prec-dscp dscp1...dscp8** global configuration command, the settings only affect packets on ingress interfaces that are configured to trust the IP precedence value. In a policy map, if you set the packet IP precedence value to a new value by using the **set ip precedence new-precedence** policy-map class configuration command, the egress DSCP value is not affected by the IP-precedence-to-DSCP map. If you want the egress DSCP value to be different than the ingress value, use the **set dscp new-dscp** policy-map class configuration command.
- If you enter or have used the **set ip dscp** command, the switch changes this command to **set dscp** in its configuration.
- You can use the **set ip precedence** or the **set precedence** policy-map class configuration command to change the packet IP precedence value. This setting appears as **set ip precedence** in the switch configuration.
- A policy-map and a port trust state can both run on a physical interface. The policy-map is applied before the port trust state.

- When you configure a default traffic class by using the **class class-default** policy-map configuration command, unclassified traffic (traffic that does not meet the match criteria specified in the traffic classes) is treated as the default traffic class (**class-default**). Class default is not supported on Catalyst 2960-S switches.

Beginning in privileged EXEC mode, follow these steps to create a policy map:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	class-map [match-all match-any] <i>class-map-name</i>	<p>Create a class map, and enter class-map configuration mode.</p> <p>By default, no class maps are defined.</p> <ul style="list-style-type: none"> (Optional) Use the match-all keyword to perform a logical-AND of all matching statements under this class map. All match 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 match criteria must be matched. For <i>class-map-name</i>, specify the name of the class map. <p>If neither the match-all or match-any keyword is specified, the default is match-all.</p> <p>Note Because only one match command per class map is supported, the match-all and match-any keywords function the same.</p>
Step 3	policy-map <i>policy-map-name</i>	<p>Create a policy map by entering the policy map name, and enter policy-map configuration mode.</p> <p>By default, no policy maps are defined.</p> <p>The default behavior of a policy map is to set the DSCP to 0 if the packet is an IP packet and to set the CoS to 0 if the packet is tagged. No policing is performed.</p>
Step 4	class [<i>class-map-name</i> class-default]	<p>Defines a traffic classification, and enter policy-map class configuration mode.</p> <p>By default, no policy map class-maps are defined.</p> <p>If a traffic class has already been defined by using the class-map global configuration command, specify its name for <i>class-map-name</i> in this command.</p> <p>A class-default traffic class is pre-defined and can be added to any policy. It is always placed at the end of a policy map. With an implied match any included in the class-default class, all packets that have not already matched the other traffic classes will match class-default.</p>

	Command	Purpose
Step 5	trust [cos dscp ip-precedence]	<p>Configure the trust state, which QoS uses to generate a CoS-based or DSCP-based QoS label.</p> <p>Note This command is mutually exclusive with the set command within the same policy map. If you enter the trust command, go to Step 6.</p> <p>By default, the port is not trusted. If no keyword is specified when the command is entered, the default is dscp.</p> <p>The keywords have these meanings:</p> <ul style="list-style-type: none"> • 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. <p>For more information, see the “Configuring the CoS-to-DSCP Map” section on page 1-63.</p>
Step 6	set {dscp <i>new-dscp</i> ip precedence <i>new-precedence</i>}	<p>Classify IP traffic by setting a new value in the packet.</p> <ul style="list-style-type: none"> • For 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.

	Command	Purpose
Step 7	police <i>rate-bps burst-byte</i> [exceed-action { drop policed-dscp-transmit }]	<p>Define a policer for the classified traffic.</p> <p>By default, no policer is defined. For information on the number of policers supported, see the “Standard QoS Configuration Guidelines” section on page 1-40.</p> <ul style="list-style-type: none"> For <i>rate-bps</i>, specify average traffic rate in bits per second (b/s). The range is 8000 to 10000000000. <p>For <i>burst-byte</i>, specify the normal burst size in bytes. The range is 8000 to 1000000. On Catalyst 2960-S switches, although you can configure a rate of 8000, the minimum rate granularity is actually 16000.</p> <ul style="list-style-type: none"> (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 (by using the policed-DSCP map) and to send the packet. For more information, see the “Configuring the Policed-DSCP Map” section on page 1-65.
Step 8	exit	Return to policy map configuration mode.
Step 9	exit	Return to global configuration mode.
Step 10	interface <i>interface-id</i>	<p>Specify the port to attach to the policy map, and enter interface configuration mode.</p> <p>Valid interfaces include physical ports.</p>
Step 11	service-policy input <i>policy-map-name</i>	<p>Specify the policy-map name, and apply it to an ingress port.</p> <p>Only one policy map per ingress port is supported.</p>
Step 12	end	Return to privileged EXEC mode.
Step 13	show policy-map [<i>policy-map-name</i> [class <i>class-map-name</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 map, use the **no class** *class-map-name* policy-map configuration command. To return to the untrusted state, use the **no trust** policy-map configuration command. To remove an assigned DSCP or IP precedence value, use the **no set** {**dscp** *new-dscp* | **ip precedence** *new-precedence*} policy-map configuration command. To remove an existing policer, use the **no police** *rate-bps burst-byte* [**exceed-action** {**drop** | **policed-dscp-transmit**}] policy-map configuration command. To remove the policy map and port association, use the **no service-policy input** *policy-map-name* interface configuration command.

This example shows how to create a policy map and attach it to an ingress 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 incoming packet is trusted. If the matched traffic exceeds an average traffic rate of 48000 b/s 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 1000000 8000 exceed-action policed-dscp-transmit
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# service-policy input flowlt
```

This example shows how to create a Layer 2 MAC ACL with two permit statements and attach it to an ingress 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 dscp 63
Switch(config-pmap-c)# exit
Switch(config-pmap)# class macclass2 maclist2
Switch(config-pmap-c)# set dscp 45
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# mls qos trust cos
Switch(config-if)# service-policy input macpolicy1
```

This example shows how to create a class map that applies to both IPv4 and IPv6 traffic with the default class applied to unclassified traffic:

```
Switch(config)# ip access-list 101 permit ip any any
Switch(config)# ipv6 access-list ipv6-any permit ip any any
Switch(config)# class-map cm-1
Switch(config-cmap)# match access-group 101
Switch(config-cmap)# exit
Switch(config)# class-map cm-2
Switch(config-cmap)# match access-group name ipv6-any
Switch(config-cmap)# exit
Switch(config)# policy-map pml
Switch(config-pmap)# class cm-1
Switch(config-pmap-c)# set dscp 4
Switch(config-pmap-c)# exit
Switch(config-pmap)# class cm-2
Switch(config-pmap-c)# set dscp 6
Switch(config-pmap-c)# exit
Switch(config-pmap)# class class-default
Switch(config-pmap-c)# set dscp 10
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface G0/1
Switch(config-if)# switch mode access
Switch(config-if)# service-policy input pml
```

This example shows that when a child-level policy map is attached below a class, an action must be specified for the class:

```
Switch(config)# policy-map vlan-plcmap
Switch(config-pmap)# class cm-5
Switch(config-pmap-c)# set dscp 7
Switch(config-pmap-c)# service-policy port-plcmap-1
```

This example shows how to configure default traffic class to a policy map:

```
Switch# configure terminal
Switch(config)# class-map cm-3
Switch(config-cmap)# match ip dscp 30
Switch(config-cmap)# match protocol ipv6
Switch(config-cmap)# exit
Switch(config)# class-map cm-4
Switch(config-cmap)# match ip dscp 40
Switch(config-cmap)# match protocol ip
Switch(config-cmap)# exit
Switch(config)# policy-map pm3
Switch(config-pmap)# class class-default
Switch(config-pmap)# set dscp 10
Switch(config-pmap-c)# exit
Switch(config-pmap)# class cm-3
Switch(config-pmap-c) set dscp 4
Switch(config-pmap-c)# exit
Switch(config-pmap)# class cm-4
Switch(config-pmap-c)# trust cos
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
```

This example shows how the default traffic class is automatically placed at the end of policy-map pm3 even though class-default was configured first:

```
Switch# show policy-map pm3
Policy Map pm3
  Class cm-3
    set dscp 4
  Class cm-4
    trust cos
  Class class-default
    police 8000 80000 exceed-action drop
Switch#
```

Classifying, Policing, and Marking Traffic by Using Aggregate Policers



Note

To use policing and marking, the switch must be running the LAN Base image.

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.

Beginning in privileged EXEC mode, follow these steps to create an aggregate policer:

	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 }	<p>Define the policer parameters that can be applied to multiple traffic classes within the same policy map.</p> <p>By default, no aggregate policer is defined. For information on the number of policers supported, see the “Standard QoS Configuration Guidelines” section on page 1-40.</p> <ul style="list-style-type: none"> For <i>aggregate-policer-name</i>, specify the name of the aggregate policer. <p>For <i>rate-bps</i>, specify average traffic rate in bits per second (b/s). The range is 8000 to 10000000000. (On Catalyst 2960-S switches, although you can configure a rate of 8000, the minimum rate granularity is actually 16000.)</p> <ul style="list-style-type: none"> 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 (by using the policed-DSCP map) and to send the packet. For more information, see the “Configuring the Policed-DSCP Map” section on page 1-65.
Step 3	class-map [match-all match-any] <i>class-map-name</i>	Create a class map to classify traffic as necessary. For more information, see the “Classifying Traffic by Using Class Maps” section on page 1-53.
Step 4	policy-map <i>policy-map-name</i>	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 Traffic on Physical Ports by Using Policy Maps” section on page 1-55.
Step 5	class [<i>class-map-name</i> class-default]	Define a traffic classification, and enter policy-map class configuration mode. The class-default keyword is not supported on Catalyst 2960-S switches. For more information, see the “Classifying, Policing, and Marking Traffic on Physical Ports by Using Policy Maps” section on page 1-55.
Step 6	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 7	exit	Return to global configuration mode.
Step 8	interface <i>interface-id</i>	Specify the port to attach to the policy map, and enter interface configuration mode. Valid interfaces include physical ports.

	Command	Purpose
Step 9	service-policy input <i>policy-map-name</i>	Specify the policy-map name, and apply it to an ingress port. Only one policy map per ingress port is supported.
Step 10	end	Return to privileged EXEC mode.
Step 11	show mls qos aggregate-policer <i>[aggregate-policer-name]</i>	Verify your entries.
Step 12	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 configuration mode. 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 incoming 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 b/s and a normal burst size of 8000 bytes, its DSCP is marked down (based on the policed-DSCP map) and sent. The policy map is attached to an ingress 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 dscp 56
Switch(config-pmap-c)# police aggregate transmit1
Switch(config-pmap-c)# exit
Switch(config-pmap)# class class-default
Switch(config-pmap-c)# set dscp 10
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# service-policy input aggflow1
Switch(config-if)# exit
```

Configuring DSCP Maps

These sections contain this configuration information:

- [Configuring the CoS-to-DSCP Map, page 1-63](#) (optional)
- [Configuring the IP-Precedence-to-DSCP Map, page 1-64](#) (optional)
- [Configuring the Policed-DSCP Map, page 1-65](#) (optional, unless the null settings in the map are not appropriate)
- [Configuring the DSCP-to-CoS Map, page 1-66](#) (optional)
- [Configuring the DSCP-to-DSCP-Mutation Map, page 1-67](#) (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 incoming packets to a DSCP value that QoS uses internally to represent the priority of the traffic.

[Table 1-13](#) shows the default CoS-to-DSCP map.

Table 1-13 Default CoS-to-DSCP Map

CoS Value	DSCP Value
0	0
1	8
2	16
3	24
4	32
5	40
6	48
7	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 <i>dscp1...dscp8</i>	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.

	Command	Purpose
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 incoming packets to a DSCP value that QoS uses internally to represent the priority of the traffic.

Table 1-14 shows the default IP-precedence-to-DSCP map:

Table 1-14 Default IP-Precedence-to-DSCP Map

IP Precedence Value	DSCP Value
0	0
1	8
2	16
3	24
4	32
5	40
6	48
7	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.

	Command	Purpose
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 a policing and marking action.

The default policed-DSCP map is a null map, which maps an incoming 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. <ul style="list-style-type: none"> 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 egress queues.

Table 1-15 shows the default DSCP-to-CoS map.

Table 1-15 Default DSCP-to-CoS Map

DSCP Value	CoS Value
0–7	0
8–15	1
16–23	2
24–31	3
32–39	4
40–47	5
48–55	6
56–63	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. <ul style="list-style-type: none"> 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 02
2 : 02 02 02 02 00 03 03 03 03 03 03
3 : 03 03 00 04 04 04 04 04 04 04 04
4 : 00 05 05 05 05 05 05 05 05 00 06
5 : 00 06 06 06 06 06 07 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

**Note**

Catalyst 2960-S switches do not support ingress queuing.

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 on an ingress port. The default DSCP-to-DSCP-mutation map is a null map, which maps an incoming 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. <ul style="list-style-type: none"> 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.

	Command	Purpose
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 gigabitethernet0/1
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 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

**Note**

Catalyst 2960-S switches do not support ingress queueing.

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 contain this configuration information:

- [Mapping DSCP or CoS Values to an Ingress Queue and Setting WTD Thresholds, page 1-69](#) (optional)
- [Allocating Buffer Space Between the Ingress Queues, page 1-71](#) (optional)
- [Allocating Bandwidth Between the Ingress Queues, page 1-72](#) (optional)
- [Configuring the Ingress Priority Queue, page 1-73](#) (optional)

Mapping DSCP or CoS Values to an Ingress Queue and Setting WTD Thresholds

**Note**

Catalyst 2960-S switches do not support ingress queueing.

You can prioritize traffic by placing packets with particular DSCPs or CoSs into certain queues and 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> <i>dscp1</i>...<i>dscp8</i> or mls qos srr-queue input cos-map queue <i>queue-id</i> threshold <i>threshold-id</i> <i>cos1</i>...<i>cos8</i>	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. <ul style="list-style-type: none"> 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</i>...<i>dscp8</i>, enter up to eight values, and separate each value with a space. The range is 0 to 63. For <i>cos1</i>...<i>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. <ul style="list-style-type: none"> For <i>queue-id</i>, the range is 1 to 2. For <i>threshold-percentage1</i> <i>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



Note Catalyst 2960-S switches do not support ingress queueing.

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 incoming 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. Note Catalyst 2960-S switches do not support ingress queueing.
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



Note Catalyst 2960-S switches do not support ingress queueing.

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. 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 1-73.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos interface queueing or show mls qos input-queue	Verify your entries. Note Catalyst 2960-S switches do not support ingress queueing.
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



Note Catalyst 2960-S switches do not support ingress queueing.

You should use the 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 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 <i>queue-id</i> bandwidth <i>weight</i>	Assign a queue as the priority queue and guarantee bandwidth on the stack or 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. <ul style="list-style-type: none"> For <i>queue-id</i>, the range is 1 to 2. For bandwidth <i>weight</i>, assign the bandwidth percentage of the stack or 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. Note Catalyst 2960-S switches do not support ingress queueing.
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 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 queues be serviced and which technique (shaped, shared, or both) should be used?

These sections contain this configuration information:

- [Configuration Guidelines, page 1-74](#)
- [Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set, page 1-74 \(optional\)](#)
- [Mapping DSCP or CoS Values to an Egress Queue and to a Threshold ID, page 1-76 \(optional\)](#)
- [Configuring SRR Shaped Weights on Egress Queues, page 1-78 \(optional\)](#)
- [Configuring SRR Shared Weights on Egress Queues, page 1-79 \(optional\)](#)
- [Configuring the Egress Expedite Queue, page 1-80 \(optional\)](#)
- [Limiting the Bandwidth on an Egress Interface, page 1-80 \(optional\)](#)

Configuration Guidelines

Follow these guidelines when the expedite queue is enabled or the egress queues are serviced based on their SRR weights:

- If the egress expedite queue is enabled, it overrides the SRR shaped and shared weights for queue 1.
- If the egress expedite queue is disabled and the SRR shaped and shared weights are configured, the shaped mode overrides the shared mode for queue 1, and SRR services this queue in shaped mode.
- If the egress expedite queue is disabled and the SRR shaped weights are not configured, SRR services this queue in shared mode.

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 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 commands.

Each threshold value is a percentage of the queues allocated buffers, 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 default settings are suitable for most situations. You should change them only when you have a thorough understanding of the egress queues and if these settings do not meet your QoS solution.

Beginning in privileged EXEC mode, follow these steps to configure the memory allocation and to 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>	<p>Allocate buffers to a queue-set.</p> <p>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.</p> <ul style="list-style-type: none"> 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. For <i>allocation1</i>, <i>allocation3</i>, and <i>allocation4</i>, the range is 0 to 99. For <i>allocation2</i>, the range is 1 to 100 (including the CPU buffer). <p>Allocate buffers according to the importance of the traffic; for example, give a large percentage of the buffer to the queue with the best-effort traffic.</p>
Step 3	mls qos queue-set output <i>qset-id</i> threshold <i>queue-id drop-threshold1</i> <i>drop-threshold2 reserved-threshold</i> <i>maximum-threshold</i>	<p>Configure the WTD thresholds, guarantee the availability of buffers, and configure the maximum memory allocation for the queue-set (four egress queues per port).</p> <p>By default, the WTD thresholds for queues 1, 3, and 4 are set to 100 percent. The thresholds for queue 2 are set to 200 percent. The reserved thresholds for queues 1, 2, 3, and 4 are set to 50 percent. The maximum thresholds for all queues are set to 400 percent.</p> <ul style="list-style-type: none"> 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 3200 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 queue in the full condition 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 3200 percent.
Step 4	interface <i>interface-id</i>	Specify the port of the outbound traffic, and enter interface configuration mode.
Step 5	queue-set <i>qset-id</i>	<p>Map the port to a queue-set.</p> <p>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.</p>
Step 6	end	Return to privileged EXEC mode.

	Command	Purpose
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 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 gigabitethernet0/1
Switch(config-if)# queue-set 2
```

Mapping DSCP or CoS Values to an Egress Queue and to a Threshold ID

You can prioritize traffic by placing packets with particular DSCPs or costs of service into certain queues and adjusting the queue thresholds so that packets with lower priorities are dropped.



Note

The egress queue default settings are suitable for most situations. You should change them only when you have a thorough understanding of the egress queues and 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. 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. <ul style="list-style-type: none"> • 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.
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 Egress Queues

You can specify how much of the available bandwidth is allocated to each queue. The ratio of the weights is the ratio of frequency in which the SRR scheduler sends packets from each queue.

You can configure the egress queues for shaped or shared weights, or both. Use shaping to smooth bursty traffic or to provide a smoother output over time. For information about shaped weights, see the “[SRR Shaping and Sharing](#)” section on page 1-14. For information about shared weights, see the “[Configuring SRR Shared Weights on Egress Queues](#)” section on page 1-79.

Beginning in privileged EXEC mode, follow these steps to assign the shaped weights and to enable bandwidth shaping on the four egress queues mapped to 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 of the outbound traffic, and enter interface configuration mode.
Step 3	srr-queue bandwidth shape <i>weight1 weight2 weight3 weight4</i>	<p>Assign SRR weights to the egress queues.</p> <p>By default, <i>weight1</i> is set to 25; <i>weight2</i>, <i>weight3</i>, and <i>weight4</i> are set to 0, and these queues are in shared mode.</p> <p>For <i>weight1 weight2 weight3 weight4</i>, enter the weights to control the percentage of the port that is shaped. The inverse ratio ($1/\textit{weight}$) controls the shaping bandwidth for this queue. Separate each value with a space. The range is 0 to 65535.</p> <p>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.</p> <p>The shaped mode overrides the shared mode.</p>
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 gigabitethernet0/1
Switch(config-if)# srr-queue bandwidth shape 8 0 0 0
```

Configuring SRR Shared Weights on Egress 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 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 default settings are suitable for most situations. You should change them only when you have a thorough understanding of the egress queues and 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 the four egress queues mapped to 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 of the outbound traffic, 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 gigabitethernet0/1
Switch(config-if)# srr-queue bandwidth share 1 2 3 4
```

Configuring the Egress Expedite Queue

You can ensure that certain packets have priority over all others by queuing them in the egress expedite queue. SRR services this queue until it is empty before servicing the other queues.

Beginning in privileged EXEC mode, follow these steps to enable the egress expedite queue. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos	Enable QoS on a switch.
Step 3	interface <i>interface-id</i>	Specify the egress port, and enter interface configuration mode.
Step 4	priority-queue out	Enable the egress expedite queue, which is disabled by default. When you configure this command, the SRR weight and queue size ratios are affected because there is one less queue participating in SRR. This means that <i>weight1</i> in the srr-queue bandwidth shape or the srr-queue bandwidth share command is ignored (not used in the ratio calculation).
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 egress expedite queue, use the **no priority-queue out** interface configuration command.

This example shows how to enable the egress expedite queue when the SRR weights are configured. The egress expedite queue overrides the configured SRR weights.

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# srr-queue bandwidth shape 25 0 0 0
Switch(config-if)# srr-queue bandwidth share 30 20 25 25
Switch(config-if)# priority-queue out
Switch(config-if)# end
```

Limiting the Bandwidth on an Egress Interface

You can limit the bandwidth on an egress port. 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 default settings are suitable for most situations. You should change them only when you have a thorough understanding of the egress queues and if these settings do not meet your QoS solution.

Beginning in privileged EXEC mode, follow these steps to limit the bandwidth on an egress 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 to be rate limited, and enter interface configuration mode.

	Command	Purpose
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. By default, the port is not rate limited and is set to 100 percent.
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 port to 80 percent:

```
Switch(config)# interface gigabitethernet0/1
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 Mb/s. 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 1-16](#):

Table 1-16 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. Note Catalyst 2960-S switches do not support ingress queueing.
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 configured policers, the queueing strategy, and the ingress and egress statistics. Note Catalyst 2960-S switches do not support ingress queueing.

Table 1-16 Commands for Displaying Standard QoS Information (continued)

Command	Purpose
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 QoS settings for the egress queues.
show policy-map [<i>policy-map-name</i> [class <i>class-map-name</i>]]	Display QoS policy maps, which define classification criteria for incoming traffic. Note Do not use the show policy-map interface privileged EXEC command to display classification information for incoming traffic. The control-plane and interface keywords are not supported, and the statistics shown in the display should be ignored.
show running-config include rewrite	Display the DSCP transparency setting.



CHAPTER 1

Configuring Static IP Unicast Routing

This chapter describes how to configure IP Version 4 (IPv4) static IP unicast routing on the Catalyst 2960-S and 2960 switch. Static routing is supported only on switched virtual interfaces (SVIs) and not on physical interfaces. The switch does not support routing protocols.

Unless otherwise noted, the term *switch* refers to a standalone switch and a switch stack. A switch stack operates and appears as a single switch to the routers in the network.



Note

Stacking is supported only on Catalyst 2960-S switches.

For complete syntax and usage information for the commands used in this chapter, see the *Cisco IOS IP Command Reference, Volume 1 of 3: Addressing and Services, Release 12.2*

- [Understanding IP Routing, page 1-1](#)
- [Steps for Configuring Routing, page 1-3](#)
- [Enabling IP Unicast Routing, page 1-4](#)
- [Configuring Static Unicast Routes, page 1-5](#)
- [Monitoring and Maintaining the IP Network, page 1-5](#)



Note


When configuring routing parameters on the switch and to allocate system resources to maximize the number of unicast routes allowed, use the **sdm prefer lanbase-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 1, “Configuring SDM Templates”](#) or see 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 to route traffic between the VLANs, referred to as inter-VLAN routing. You configure one or more routers to route traffic to the appropriate destination VLAN.

[Figure 1-1](#) 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 1-1 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 uses the routing table to find 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.

When static routing is enabled on Switch A and B, the router device is no longer needed to route packets.

Types of Routing

Routers and Layer 3 switches can route packets in these ways:

- Using default routing to send traffic with a destination unknown to the router to a default outlet or destination
- Using static routes to forward packets from predetermined ports through a single path into and out of a network
- Dynamically calculating routes by using a routing protocol

The switch supports static routes and default routes, It does not support routing protocols.

IP Routing and Switch Stacks



Note

Stacking is supported only on Catalyst 2960-S switches.

A switch stack appears to the network as a single switch, regardless of which switch in the stack is connected to a peer. For additional information about switch stack operation, see [Chapter 1, “Managing Switch Stacks.”](#)

Stack master functions:

- The MAC address of the stack master is used as the router MAC address for the whole stack, and all outside devices use this address to send IP packets to the stack.
- All IP packets that require software forwarding or processing go through the CPU of the stack master.

Stack members functions:

- Act as routing standby switches, taking over if elected as the new stack master when the stack master fails.
- Program the routes into hardware.

If a stack master fails, the stack detects that the stack master is down and elects a stack member to be the new stack master. Except for a momentary interruption, the hardware continues to forward packets.

New stack master functions after election:

- Builds routing table and distributes it to stack members.
- Uses its MAC address as the router MAC address. To notify its network peers of the new MAC address, it periodically (every few seconds for 5 minutes) sends a gratuitous ARP reply with the new router MAC address.

**Note**

If you configure the persistent MAC address feature on the stack and the stack master changes, the stack MAC address does not change during the configured time period. If the previous stack master rejoins the stack as a member switch during that time period, the stack MAC address remains the MAC address of the previous stack master. See the [“Enabling Persistent MAC Address”](#) section on page 1-17.

Steps for Configuring Routing

By default, IP routing is disabled on the switch. For detailed IP routing configuration information, see the *Cisco IOS IP Configuration Guide, Release 12.2* from the Cisco.com page under **Documentation > Cisco IOS Software Releases > 12.2 Mainline > Configuration Guides**.

In these procedures, the specified interface must be 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. All Layer 3 interfaces on which routing will occur must have IP addresses assigned to them. See the [“Assigning IP Addresses to SVIs”](#) section on page 1-4.

**Note**

The switch supports 16 static routes (including user-configured routes and the default route) and any directly connected routes and default routes for the management interface. The switch can have an IP address assigned to each SVI. Before enabling routing, enter the **sdm prefer lanbase-routing** global configuration command and reload the switch.

Procedures for configuring routing:

- To support VLAN interfaces, create and configure VLANs on the switch or switch stack, and assign VLAN membership to Layer 2 interfaces. For more information, see [Chapter 1, “Configuring VLANs.”](#)
- Configure Layer 3 interfaces (SVIs).
- Enable IP routing on the switch.
- Assign IP addresses to the Layer 3 interfaces.
- Configure static routes

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, 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	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.

Use the **no ip routing** global configuration command to disable routing.

This example shows how to enable IP routing on a switch:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip routing
Switch(config)# end
```

Assigning IP Addresses to SVIs

To configure IP routing, you need to assign IP addresses to Layer 3 network interfaces. This enables communication with the hosts on those interfaces that use IP. IP routing is disabled by default, and no IP addresses are assigned to SVIs.

An IP address identifies a destination for IP packets. 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 these IP addresses.

An interface can have one primary IP address. A subnet mask identifies the bits that denote the network number in an IP address.

Beginning in privileged EXEC mode, follow these steps to assign an IP address and a network mask to an SVI

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface vlan <i>vlan_id</i>	Enter interface configuration mode, and specify the Layer 3 VLAN to configure.
Step 3	ip address <i>ip-address subnet-mask</i>	Configure the IP address and IP subnet mask.
Step 4	end	Return to privileged EXEC mode.
Step 5	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 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

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.

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.

Monitoring and Maintaining the IP Network

You can specific statistics for the routing table or database. Use the privileged EXEC commands in [Table 1-1](#) to display status:

Table 1-1 Commands to Clear IP Routes or Display Route Status

Command	Purpose
show ip route [<i>address [mask] [longer-prefixes]</i>]	Display the state of the routing table.
show ip route summary	Display the state of the routing table in summary form.
show platform ip unicast	Display platform-dependent IP unicast information.



CHAPTER 1

Configuring IPv6 Host Functions

This chapter describes how to configure IPv6 host functions on the Catalyst 2960 and 2960-S switches. For information about configuring IPv6 Multicast Listener Discovery (MLD) snooping, see [Chapter 1, “Configuring IPv6 MLD Snooping.”](#)

To enable dual stack environments (supporting both IPv4 and IPv6) on a Catalyst 2960 switch, you must configure the switch to use the a dual IPv4 and IPv6 switch database management (SDM) template. See the [“Dual IPv4 and IPv6 Protocol Stacks” section on page 1-4](#). This template is not supported on Catalyst 2960-S switches.



Note

For complete syntax and usage information for the commands used in this chapter, see the Cisco IOS documentation referenced in the procedures.

- [“Understanding IPv6” section on page 1-1](#)
- [“Configuring IPv6” section on page 1-6](#)
- [“Displaying IPv6” section on page 1-11](#)

Understanding IPv6

IPv4 users can move to IPv6 and receive services such as end-to-end security, quality of service (QoS), and globally unique addresses. The IPv6 address space reduces the need for private addresses and Network Address Translation (NAT) processing by border routers at network edges.

For information about how Cisco Systems implements IPv6, go to:

http://www.cisco.com/en/US/products/ps6553/products_ios_technology_home.html

For information about IPv6 and other features in this chapter:

- See the *Cisco IOS IPv6 Configuration Library*:
http://www.cisco.com/en/US/docs/ios/12_2t/ipv6/ipv6_vgf.html
- Use the Search field to locate the Cisco IOS software documentation. For example, if you want information about static routes, you can enter *Implementing Static Routes for IPv6* in the search field to get this document about static routes:
http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-stat_routes_ps6441_TSD_Products_Configuration_Guide_Chapter.html

These sections describe IPv6 implementation on the switch.

- [IPv6 Addresses, page 1-2](#)

- [Supported IPv6 Host Features, page 1-2](#)

IPv6 Addresses

The switch supports only IPv6 unicast addresses. It does not support site-local unicast addresses, anycast addresses, or multicast addresses.

The IPv6 128-bit addresses are represented as a series of eight 16-bit hexadecimal fields separated by colons in the format: n:n:n:n:n:n:n:n. This is an example of an IPv6 address:

```
2031:0000:130F:0000:0000:09C0:080F:130B
```

For easier implementation, leading zeros in each field are optional. This is the same address without leading zeros:

```
2031:0:130F:0:0:9C0:80F:130B
```

You can also use two colons (::) to represent successive hexadecimal fields of zeros, but you can use this short version only once in each address:

```
2031:0:130F::09C0:080F:130B
```

For more information about IPv6 address formats, address types, and the IPv6 packet header, see the “Implementing IPv6 Addressing and Basic Connectivity” chapter of *Cisco IOS IPv6 Configuration Library* on Cisco.com.

In the “Implementing Addressing and Basic Connectivity” chapter, these sections apply to the Catalyst 2960 switch:

- IPv6 Address Formats
- IPv6 Address Output Display
- Simplified IPv6 Packet Header

Supported IPv6 Host Features

These sections describe the IPv6 protocol features supported by the switch:

- [128-Bit Wide Unicast Addresses, page 1-3](#)
- [DNS for IPv6, page 1-3](#)
- [ICMPv6, page 1-3](#)
- [Neighbor Discovery, page 1-3](#)
- [IPv6 Stateless Autoconfiguration and Duplicate Address Detection, page 1-4](#)
- [IPv6 Applications, page 1-4](#)
- [Dual IPv4 and IPv6 Protocol Stacks, page 1-4](#)
- [SNMP and Syslog Over IPv6, page 1-5](#)
- [HTTP\(S\) Over IPv6, page 1-5](#)

Support on the switch includes expanded address capability, header format simplification, improved support of extensions and options, and hardware parsing of the extension header. The switch supports hop-by-hop extension header packets, which are routed or bridged in software.

128-Bit Wide Unicast Addresses

The switch supports aggregatable global unicast addresses and link-local unicast addresses. It does not support site-local unicast addresses.

- Aggregatable global unicast addresses are IPv6 addresses from the aggregatable global unicast prefix. The address structure enables strict aggregation of routing prefixes and limits the number of routing table entries in the global routing table. These addresses are used on links that are aggregated through organizations and eventually to the Internet service provider.

These addresses are defined by a global routing prefix, a subnet ID, and an interface ID. Current global unicast address allocation uses the range of addresses that start with binary value 001 (2000::/3). Addresses with a prefix of 2000::/3(001) through E000::/3(111) must have 64-bit interface identifiers in the extended unique identifier (EUI)-64 format.

- Link local unicast addresses can be automatically configured on any interface by using the link-local prefix FE80::/10(1111 1110 10) and the interface identifier in the modified EUI format. Link-local addresses are used in the neighbor discovery protocol (NDP) and the stateless autoconfiguration process. Nodes on a local link use link-local addresses and do not require globally unique addresses to communicate. IPv6 routers do not forward packets with link-local source or destination addresses to other links.

For more information, see the section about IPv6 unicast addresses in the “Implementing IPv6 Addressing and Basic Connectivity” chapter in the *Cisco IOS IPv6 Configuration Library* on Cisco.com.

DNS for IPv6

IPv6 supports Domain Name System (DNS) record types in the DNS name-to-address and address-to-name lookup processes. The DNS AAAA resource record types support IPv6 addresses and are equivalent to an A address record in IPv4. The switch supports DNS resolution for IPv4 and IPv6.

ICMPv6

The Internet Control Message Protocol (ICMP) in IPv6 generates error messages, such as ICMP destination unreachable messages, to report errors during processing and other diagnostic functions. In IPv6, ICMP packets are also used in the neighbor discovery protocol and path MTU discovery.

Neighbor Discovery

The switch supports NDP for IPv6, a protocol running on top of ICMPv6, and static neighbor entries for IPv6 stations that do not support NDP. The IPv6 neighbor discovery process uses ICMP messages and solicited-node multicast addresses to determine the link-layer address of a neighbor on the same network (local link), to verify the reachability of the neighbor, and to keep track of neighboring routers.

The switch supports ICMPv6 redirect for routes with mask lengths less than 64 bits. ICMP redirect is not supported for host routes or for summarized routes with mask lengths greater than 64 bits.

Neighbor discovery throttling ensures that the switch CPU is not unnecessarily burdened while it is in the process of obtaining the next hop forwarding information to route an IPv6 packet. The switch drops any additional IPv6 packets whose next hop is the same neighbor that the switch is actively trying to resolve. This drop avoids further load on the CPU.

IPv6 Stateless Autoconfiguration and Duplicate Address Detection

The switch uses stateless autoconfiguration to manage link, subnet, and site addressing changes, such as management of host and mobile IP addresses. A host autonomously configures its own link-local address, and booting nodes send router solicitations to request router advertisements for configuring interfaces.

For more information about autoconfiguration and duplicate address detection, see the “Implementing IPv6 Addressing and Basic Connectivity” chapter of *Cisco IOS IPv6 Configuration Library* on Cisco.com.

IPv6 Applications

The switch has IPv6 support for these applications:

- Ping, traceroute, and Telnet
- Secure Shell (SSH) over an IPv6 transport
- HTTP server access over IPv6 transport
- DNS resolver for AAAA over IPv6 transport
- Cisco Discovery Protocol (CDP) support for IPv6 addresses


For more information about managing these applications, see the “Managing Cisco IOS Applications over IPv6” chapter and the “Implementing IPv6 Addressing and Basic Connectivity” chapter in the *Cisco IOS IPv6 Configuration Library* on Cisco.com.

Dual IPv4 and IPv6 Protocol Stacks

On a Catalyst 2960 switch, you must use the dual IPv4 and IPv6 template to allocate ternary content addressable memory (TCAM) usage to both IPv4 and IPv6 protocols.

Figure 1-1 shows a router forwarding both IPv4 and IPv6 traffic through the same interface, based on the IP packet and destination addresses.

Figure 1-1 Dual IPv4 and IPv6 Support on an Interface



Use the dual IPv4 and IPv6 switch database management (SDM) template on a Catalyst 2960 switch to enable dual stack environments (supporting both IPv4 and IPv6). For more information about the dual IPv4 and IPv6 SDM template, see Chapter 1, “Configuring SDM Templates.”

The dual IPv4 and IPv6 templates on Catalyst 2960 switches allow the switch to be used in dual stack environments.

- If you try to configure IPv6 without first selecting a dual IPv4 and IPv6 template, a warning message appears.
- In IPv4-only environments, the switch applies IPv4 QoS and ACLs in hardware. IPv6 packets are not supported.
- In dual IPv4 and IPv6 environments, the switch applies IPv4 QoS and ACLs in hardware.
- IPv6 QoS and ACLs are not supported.
- If you do not plan to use IPv6, do not use the dual stack template because this template results in less TCAM capacity for each resource.

For more information about IPv4 and IPv6 protocol stacks, see the “Implementing IPv6 Addressing and Basic Connectivity” chapter of *Cisco IOS IPv6 Configuration Library* on Cisco.com.

SNMP and Syslog Over IPv6

To support both IPv4 and IPv6, IPv6 network management requires both IPv6 and IPv4 transports. Syslog over IPv6 supports address data types for these transports.

SNMP and syslog over IPv6 provide these features:

- Support for both IPv4 and IPv6
- IPv6 transport for SNMP and to modify the SNMP agent to support traps for an IPv6 host
- SNMP- and syslog-related MIBs to support IPv6 addressing
- Configuration of IPv6 hosts as trap receivers

For support over IPv6, SNMP modifies the existing IP transport mapping to simultaneously support IPv4 and IPv6. These SNMP actions support IPv6 transport management:

- Opens User Datagram Protocol (UDP) SNMP socket with default settings
- Provides a new transport mechanism called *SR_IPV6_TRANSPORT*
- Sends SNMP notifications over IPv6 transport
- Supports SNMP-named access lists for IPv6 transport
- Supports SNMP proxy forwarding using IPv6 transport
- Verifies SNMP Manager feature works with IPv6 transport

For information on SNMP over IPv6, including configuration procedures, see the “Managing Cisco IOS Applications over IPv6” chapter in the *Cisco IOS IPv6 Configuration Library* on Cisco.com.

For information about syslog over IPv6, including configuration procedures, see the “Implementing IPv6 Addressing and Basic Connectivity” chapter in the *Cisco IOS IPv6 Configuration Library* on Cisco.com.

HTTP(S) Over IPv6

The HTTP client sends requests to both IPv4 and IPv6 HTTP servers, which respond to requests from both IPv4 and IPv6 HTTP clients. URLs with literal IPv6 addresses must be specified in hexadecimal using 16-bit values between colons.

The accept socket call chooses an IPv4 or IPv6 address family. The accept socket is either an IPv4 or IPv6 socket. The listening socket continues to listen for both IPv4 and IPv6 signals that indicate a connection. The IPv6 listening socket is bound to an IPv6 wildcard address.

The underlying TCP/IP stack supports a dual-stack environment. HTTP relies on the TCP/IP stack and the sockets for processing network-layer interactions.

Basic network connectivity (**ping**) must exist between the client and the server hosts before HTTP connections can be made.

For more information, see the “Managing Cisco IOS Applications over IPv6” chapter in the *Cisco IOS IPv6 Configuration Library* on Cisco.com.

IPv6 and Switch Stacks

The switch supports IPv6 forwarding across the stack and IPv6 host functionality on the stack master. The stack master runs IPv6 host functionality and IPv6 applications.

While the new stack master is being elected and is resetting, the switch stack does not forward IPv6 packets. The stack MAC address changes, which also changes the IPv6 address. When you specify the stack IPv6 address with an extended unique identifier (EUI) by using the **ipv6 address ipv6-prefix/prefix length eui-64** interface configuration command, the address is based on the interface MAC address. See the “[Configuring IPv6 Addressing and Enabling IPv6 Host](#)” section on page 1-7.

If you configure the persistent MAC address feature on the stack and the stack master changes, the stack MAC address does not change for approximately 4 minutes. For more information, see the “[Enabling Persistent MAC Address](#)” section on page 1-17 in Chapter 1, “Managing Switch Stacks.”

Configuring IPv6

These sections contain this IPv6 forwarding configuration information:

- [Default IPv6 Configuration](#), page 1-6
- [Configuring IPv6 Addressing and Enabling IPv6 Host](#), page 1-7
- [Configuring IPv6 ICMP Rate Limiting](#), page 1-9
- [Configuring Static Routes for IPv6](#), page 1-10

Default IPv6 Configuration

Table 1-1 shows the default IPv6 configuration.

Table 1-1 Default IPv6 Configuration

Feature	Default Setting
SDM template	Default
IPv6 addresses	None configured.

Configuring IPv6 Addressing and Enabling IPv6 Host

This section describes how to assign IPv6 addresses to individual Layer 3 interfaces and to globally forward IPv6 traffic on the switch.

Before configuring IPv6 on the switch, consider these guidelines:

- Be sure to select a dual IPv4 and IPv6 SDM template.
- In the **ipv6 address** interface configuration command, you must enter the *ipv6-address* and *ipv6-prefix* variables with the address specified in hexadecimal using 16-bit values between colons. The *prefix-length* variable (preceded by a slash [/]) is a decimal value that shows how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address).

To forward IPv6 traffic on an interface, you must configure a global IPv6 address on that interface. Configuring an IPv6 address on an interface automatically configures a link-local address and activates IPv6 for the interface. The configured interface automatically joins these required multicast groups for that link:

- solicited-node multicast group FF02:0:0:0:1:ff00::/104 for each unicast address assigned to the interface (this address is used in the neighbor discovery process.)
- all-nodes link-local multicast group FF02::1
- all-routers link-local multicast group FF02::2

For more information about configuring IPv6, see the “Implementing Addressing and Basic Connectivity for IPv6” chapter in the *Cisco IOS IPv6 Configuration Library* on Cisco.com.

Beginning in privileged EXEC mode, follow these steps to assign an IPv6 address to a Layer 3 interface and enable IPv6 forwarding:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	sdm prefer dual-ipv4-and-ipv6 default	Select the SDM template that supports IPv4 and IPv6.
Step 3	end	Return to privileged EXEC mode.
Step 4	reload	Reload the operating system.
Step 5	configure terminal	Enter global configuration mode after the switch reloads.
Step 6	interface interface-id	Enter interface configuration mode, and specify the interface to configure.

	Command	Purpose
Step 7	ipv6 address <i>ipv6-prefix/prefix length eui-64</i>	Specify a global IPv6 address with an extended unique identifier (EUI) in the low-order 64 bits of the IPv6 address. Specify only the network prefix; the last 64 bits are automatically computed from the switch MAC address. This enables IPv6 processing on the interface.
	or	
	ipv6 address <i>ipv6-address/prefix length</i>	Manually configure an IPv6 address on the interface.
	or	
	ipv6 address <i>ipv6-address link-local</i>	Specify a link-local address on the interface to be used instead of the link-local address that is automatically configured when IPv6 is enabled on the interface. This command enables IPv6 processing on the interface.
	or	
	ipv6 enable	Automatically configure an IPv6 link-local address on the interface, and enable the interface for IPv6 processing. The link-local address can only be used to communicate with nodes on the same link.
Step 8	exit	Return to global configuration mode.
Step 9	end	Return to privileged EXEC mode.
Step 10	show ipv6 interface <i>interface-id</i>	Verify your entries.
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove an IPv6 address from an interface, use the **no ipv6 address** *ipv6-prefix/prefix length eui-64* or **no ipv6 address** *ipv6-address link-local* interface configuration command. To remove all manually configured IPv6 addresses from an interface, use the **no ipv6 address** interface configuration command without arguments. To disable IPv6 processing on an interface that has not been explicitly configured with an IPv6 address, use the **no ipv6 enable** interface configuration command. To globally disable IPv6 routing, use the **no ipv6 unicast-routing** global configuration command.

This example shows how to enable IPv6 with both a link-local address and a global address based on the IPv6 prefix 2001:0DB8:c18:1::/64. The EUI-64 interface ID is used in the low-order 64 bits of both addresses. Output from the **show ipv6 interface EXEC** command shows how the interface ID (20B:46FF:FE2F:D940) is appended to the link-local prefix FE80::/64 of the interface.

```
Switch(config)# sdm prefer dual-ipv4-and-ipv6 default
Switch(config)# interface gigabitEthernet0/1
Switch(config-if)# ipv6 address 2001:0DB8:c18:1::/64 eui 64
Switch(config-if)# end
Switch# show ipv6 interface gigabitEthernet0/1
GigabitEthernet0/1 is up, line protocol is up
IPv6 is enabled, link-local address is FE80::20B:46FF:FE2F:D940
Global unicast address(es):
  2001:0DB8:c18:1:20B:46FF:FE2F:D940, subnet is 2001:0DB8:c18:1::/64 [EUI]
Joined group address(es):
  FF02::1
  FF02::2
  FF02::1:FF2F:D940
MTU is 1500 bytes
ICMP error messages limited to one every 100 milliseconds
ICMP redirects are enabled
ND DAD is enabled, number of DAD attempts: 1
ND reachable time is 30000 milliseconds
ND advertised reachable time is 0 milliseconds
ND advertised retransmit interval is 0 milliseconds
```

```

ND router advertisements are sent every 200 seconds
ND router advertisements live for 1800 seconds
Hosts use stateless autoconfig for addresses.
0/1

```

Configuring IPv6 ICMP Rate Limiting

ICMP rate limiting is enabled by default with a default interval between error messages of 100 milliseconds and a bucket size (maximum number of tokens to be stored in a bucket) of 10.

Beginning in privileged EXEC mode, follow these steps to change the ICMP rate-limiting parameters:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ipv6 icmp error-interval <i>interval</i> [<i>bucketsize</i>]	Configure the interval and bucket size for IPv6 ICMP error messages: <ul style="list-style-type: none"> <i>interval</i>—The interval (in milliseconds) between tokens being added to the bucket. The range is from 0 to 2147483647 milliseconds. <i>bucketsize</i>—(Optional) The maximum number of tokens stored in the bucket. The range is from 1 to 200.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ipv6 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 to the default configuration, use the **no ipv6 icmp error-interval** global configuration command.

This example shows how to configure an IPv6 ICMP error message interval of 50 milliseconds and a bucket size of 20 tokens.

```
Switch(config)#ipv6 icmp error-interval 50 20
```

Configuring Static Routes for IPv6

Beginning in privileged EXEC mode, follow these steps to configure an IPv6 static route:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ipv6 route <i>ipv6-prefix/prefix length</i> { <i>ipv6-address</i> <i>interface-id</i> [<i>ipv6-address</i>]} [<i>administrative distance</i>]	<p>Configure a static IPv6 route.</p> <ul style="list-style-type: none"> • <i>ipv6-prefix</i>—The IPv6 network that is the destination of the static route. It can also be a hostname when static host routes are configured. • <i>/prefix length</i>—The length of the IPv6 prefix. A decimal value that shows how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address). A slash mark must precede the decimal value. • <i>ipv6-address</i>—The IPv6 address of the next hop that can be used to reach the specified network. The IPv6 address of the next hop need not be directly connected; recursion is done to find the IPv6 address of the directly connected next hop. The address must be specified in hexadecimal using 16-bit values between colons. • <i>interface-id</i>—Specify direct static routes from point-to-point and broadcast interfaces. With point-to-point interfaces, there is no need to specify the IPv6 address of the next hop. With broadcast interfaces, you should always specify the IPv6 address of the next hop, or ensure that the specified prefix is assigned to the link, specifying a link-local address as the next hop. You can optionally specify the IPv6 address of the next hop to which packets are sent. <p>Note You must specify an <i>interface-id</i> when using a link-local address as the next hop (the link-local next hop must also be an adjacent router).</p> <ul style="list-style-type: none"> • <i>administrative distance</i>—(Optional) An administrative distance. The range is 1 to 254; the default value is 1, which gives static routes precedence over any other type of route except connected routes. To configure a floating static route, use an administrative distance greater than that of the dynamic routing protocol.
Step 3	end	Return to privileged EXEC mode.

	Command	Purpose
Step 4	<p>show ipv6 static [<i>ipv6-address</i> <i>ipv6-prefix/prefix length</i>] [interface <i>interface-id</i>] [recursive] [detail]</p> <p>or</p> <p>show ipv6 route static [<i>updated</i>]</p>	<p>Verify your entries by displaying the contents of the IPv6 routing table.</p> <ul style="list-style-type: none"> • interface <i>interface-id</i>—(Optional) Display only those static routes with the specified interface as an egress interface. • recursive—(Optional) Display only recursive static routes. The recursive keyword is mutually exclusive with the interface keyword, but it can be used with or without the IPv6 prefix included in the command syntax. • detail—(Optional) Display this additional information: <ul style="list-style-type: none"> – For valid recursive routes, the output path set, and maximum resolution depth. – For invalid routes, the reason why the route is not valid.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a configured static route, use the **no ipv6 route** *ipv6-prefix/prefix length* {*ipv6-address* | *interface-id* [*ipv6-address*]} [*administrative distance*] global configuration command.

This example shows how to configure a floating static route with an administrative distance of 130 to an interface:

```
Switch(config)# ipv6 route 2001:0DB8::/32 gigabitethernet0/1 130
```

For more information about configuring static IPv6 routing, see the “Implementing Static Routes for IPv6” chapter in the *Cisco IOS IPv6 Configuration Library* on Cisco.com.

Displaying IPv6

For complete syntax and usage information on these commands, see the Cisco IOS command reference publications.

Table 1-2 shows the privileged EXEC commands for monitoring IPv6 on the switch.

Table 1-2 Commands for Monitoring IPv6

Command	Purpose
show ipv6 access-list	Display a summary of access lists.
show ipv6 interface <i>interface-id</i>	Display IPv6 interface status and configuration.
show ipv6 mtu	Display IPv6 MTU per destination cache.
show ipv6 neighbors	Display IPv6 neighbor cache entries.
show ipv6 prefix-list	Display a list of IPv6 prefix lists.
show ipv6 protocols	Display IPv6 routing protocols on the switch.
show ipv6 route	Display the IPv6 route table entries.
show ipv6 static	Display IPv6 static routes.
show ipv6 traffic	Display IPv6 traffic statistics.

Table 1-3 shows the privileged EXEC commands for displaying information about IPv4 and IPv6 address types.

Table 1-3 Commands for Displaying IPv4 and IPv6 Address Types

Command	Purpose
show ip http server history	Display the previous 20 connections to the HTTP server, including the IP address accessed and the time when the connection was closed.
show ip http server connection	Display the current connections to the HTTP server, including the local and remote IP addresses being accessed.
show ip http client connection	Display the configuration values for HTTP client connections to HTTP servers.
show ip http client history	Display a list of the last 20 requests made by the HTTP client to the server.

This is an example of the output from the **show ipv6 interface** privileged EXEC command:

```
Switch# show ipv6 interface
Vlan1 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::20B:46FF:FE2F:D940
  Global unicast address(es):
    3FFE:C000:0:1:20B:46FF:FE2F:D940, subnet is 3FFE:C000:0:1::/64 [EUI]
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF2F:D940
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
<output truncated>
```

This is an example of the output from the **show ipv6 protocols** privileged EXEC command:

```
Switch# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "static"
IPv6 Routing Protocol is "rip fer"
  Interfaces:
    Vlan6
    FastEthernet0/4
    FastEthernet0/11
    FastEthernet0/12
  Redistribution:
    None
```

This is an example of the output from the **show ipv6 static** privileged EXEC command:

```
Switch# show ipv6 static
IPv6 Static routes
Code: * - installed in RIB
* ::/0 via nexthop 3FFE:C000:0:7::777, distance 1
```


This is an example of the output from the **show ipv6 neighbor** privileged EXEC command:

```
Switch# show ipv6 neighbors
IPv6 Address                               Age Link-layer Addr State Interface
3FFE:C000:0:7::777                         - 0007.0007.0007 REACH V17
3FFE:C101:113:1::33                       - 0000.0000.0033 REACH Fa1/0/13
```

This is an example of the output from the **show ipv6 route** privileged EXEC command:

```
Switch# show ipv6 route
IPv6 Routing Table - Default - 1 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
L   FF00::/8 [0/0]
    via Null0, receive
```

This is an example of the output from the **show ipv6 traffic** privileged EXEC command.

```
Switch# show ipv6 traffic
IPv6 statistics:
  Rcvd:  1 total, 1 local destination
         0 source-routed, 0 truncated
         0 format errors, 0 hop count exceeded
         0 bad header, 0 unknown option, 0 bad source
         0 unknown protocol, 0 not a router
         0 fragments, 0 total reassembled
         0 reassembly timeouts, 0 reassembly failures
  Sent:  36861 generated, 0 forwarded
         0 fragmented into 0 fragments, 0 failed
         0 encapsulation failed, 0 no route, 0 too big
         0 RPF drops, 0 RPF suppressed drops
  Mcast: 1 received, 36861 sent

ICMP statistics:
  Rcvd:  1 input, 0 checksum errors, 0 too short
         0 unknown info type, 0 unknown error type
  unreach: 0 routing, 0 admin, 0 neighbor, 0 address, 0 port
  parameter: 0 error, 0 header, 0 option
         0 hopcount expired, 0 reassembly timeout, 0 too big
         0 echo request, 0 echo reply
         0 group query, 0 group report, 0 group reduce
         1 router solicit, 0 router advert, 0 redirects
         0 neighbor solicit, 0 neighbor advert
  Sent:  10112 output, 0 rate-limited
         unreach: 0 routing, 0 admin, 0 neighbor, 0 address, 0 port
         parameter: 0 error, 0 header, 0 option
         0 hopcount expired, 0 reassembly timeout, 0 too big
         0 echo request, 0 echo reply
         0 group query, 0 group report, 0 group reduce
         0 router solicit, 9944 router advert, 0 redirects
         84 neighbor solicit, 84 neighbor advert

UDP statistics:
  Rcvd:  0 input, 0 checksum errors, 0 length errors
         0 no port, 0 dropped
  Sent:  26749 output

TCP statistics:
  Rcvd:  0 input, 0 checksum errors
  Sent:  0 output, 0 retransmitted
```




CHAPTER 1

Configuring IPv6 MLD Snooping

You can use Multicast Listener Discovery (MLD) snooping to enable efficient distribution of IP version 6 (IPv6) multicast data to clients and routers in a switched network on the Catalyst 2960 and 2960-S switches. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note Stacking is supported only on Catalyst 2960-S switches running the LAN base image.



Note To use IPv6 on a Catalyst 2960 switch, you must configure the dual IPv4 and IPv6 Switch Database Management (SDM) template on the switch. You select the template by entering the **sdm prefer dual-ipv4-and-ipv6 default** global configuration command. This template is not required on a Catalyst 2960-S switch.

For related information, see these chapters:

- For more information about SDM templates, see [Chapter 1, “Configuring SDM Templates.”](#)
- For information about IPv6 on the switch, see [Chapter 1, “Configuring IPv6 Host Functions.”](#)



Note For complete syntax and usage information for the commands used in this chapter, see the command reference for this release or the Cisco IOS documentation referenced in the procedures.

This chapter includes these sections:

- [“Understanding MLD Snooping” section on page 1-1](#)
- [“Configuring IPv6 MLD Snooping” section on page 1-5](#)
- [“Displaying MLD Snooping Information” section on page 1-11](#)

Understanding MLD Snooping

In IP version 4 (IPv4), Layer 2 switches can use Internet Group Management Protocol (IGMP) snooping to limit 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. In IPv6, MLD snooping performs a similar function. With MLD snooping, IPv6 multicast data is selectively forwarded to a list of ports that want to receive the data, instead of being flooded to all ports in a VLAN. This list is constructed by snooping IPv6 multicast control packets.

MLD is a protocol used by IPv6 multicast routers to discover the presence of multicast listeners (nodes wishing to receive IPv6 multicast packets) on its directly attached links and to discover which multicast packets are of interest to neighboring nodes. MLD is derived from IGMP; MLD version 1 (MLDv1) is equivalent to IGMPv2 and MLD version 2 (MLDv2) is equivalent to IGMPv3. MLD is a subprotocol of Internet Control Message Protocol version 6 (ICMPv6), and MLD messages are a subset of ICMPv6 messages, identified in IPv6 packets by a preceding Next Header value of 58.

The switch supports two versions of MLD snooping:

- MLDv1 snooping detects MLDv1 control packets and sets up traffic bridging based on IPv6 destination multicast addresses.
- MLDv2 basic snooping (MBSS) uses MLDv2 control packets to set up traffic forwarding based on IPv6 destination multicast addresses.

The switch can snoop on both MLDv1 and MLDv2 protocol packets and bridge IPv6 multicast data based on destination IPv6 multicast addresses.


Note

The switch does not support MLDv2 enhanced snooping (MESS), which sets up IPv6 source and destination multicast address-based forwarding.

MLD snooping can be enabled or disabled globally or per VLAN. When MLD snooping is enabled, a per-VLAN IPv6 multicast MAC address table is constructed in software and a per-VLAN IPv6 multicast address table is constructed in software and hardware. The switch then performs IPv6 multicast-address based bridging in hardware.

These sections describe some parameters of IPv6 MLD snooping:

- [MLD Messages, page 1-2](#)
- [MLD Queries, page 1-3](#)
- [Multicast Client Aging Robustness, page 1-3](#)
- [Multicast Router Discovery, page 1-3](#)
- [MLD Reports, page 1-4](#)
- [MLD Done Messages and Immediate-Leave, page 1-4](#)
- [Topology Change Notification Processing, page 1-5](#)
- [MLD Snooping in Switch Stacks, page 1-5](#)

MLD Messages

MLDv1 supports three types of messages:

- Listener Queries are the equivalent of IGMPv2 queries and are either General Queries or Multicast-Address-Specific Queries (MASQs).
- Multicast Listener Reports are the equivalent of IGMPv2 reports.
- Multicast Listener Done messages are the equivalent of IGMPv2 leave messages.

MLDv2 supports MLDv2 queries and reports, as well as MLDv1 Report and Done messages.

Message timers and state transitions resulting from messages being sent or received are the same as those of IGMPv2 messages. MLD messages that do not have valid link-local IPv6 source addresses are ignored by MLD routers and switches.

MLD Queries

The switch sends out MLD queries, constructs an IPv6 multicast address database, and generates MLD group-specific and MLD group-and-source-specific queries in response to MLD Done messages. The switch also supports report suppression, report proxying, Immediate-Leave functionality, and static IPv6 multicast MAC-address configuration.

When MLD snooping is disabled, all MLD queries are flooded in the ingress VLAN.

When MLD snooping is enabled, received MLD queries are flooded in the ingress VLAN, and a copy of the query is sent to the CPU for processing. From the received query, MLD snooping builds the IPv6 multicast address database. It detects multicast router ports, maintains timers, sets report response time, learns the querier IP source address for the VLAN, learns the querier port in the VLAN, and maintains multicast-address aging.



Note

When the IPv6 multicast router is a Catalyst 6500 switch and you are using extended VLANs (in the range 1006 to 4094), IPv6 MLD snooping must be enabled on the extended VLAN on the Catalyst 6500 switch in order for the Catalyst 2960 or 2960-S switch to receive queries on the VLAN. For normal-range VLANs (1 to 1005), it is not necessary to enable IPv6 MLD snooping on the VLAN on the Catalyst 6500 switch.

When a group exists in the MLD snooping database, the switch responds to a group-specific query by sending an MLDv1 report. When the group is unknown, the group-specific query is flooded to the ingress VLAN.

When a host wants to leave a multicast group, it can send out an MLD Done message (equivalent to IGMP Leave message). When the switch receives an MLDv1 Done message, if Immediate-Leave is not enabled, the switch sends an MASQ to the port from which the message was received to determine if other devices connected to the port should remain in the multicast group.

Multicast Client Aging Robustness

You can configure port membership removal from addresses based on the number of queries. A port is removed from membership to an address only when there are no reports to the address on the port for the configured number of queries. The default number is 2.

Multicast Router Discovery

Like IGMP snooping, MLD snooping performs multicast router discovery, with these characteristics:

- Ports configured by a user never age out.
- Dynamic port learning results from MLDv1 snooping queries and IPv6 PIMv2 packets.
- If there are multiple routers on the same Layer 2 interface, MLD snooping tracks a single multicast router on the port (the router that most recently sent a router control packet).
- Dynamic multicast router port aging is based on a default timer of 5 minutes; the multicast router is deleted from the router port list if no control packet is received on the port for 5 minutes.
- IPv6 multicast router discovery only takes place when MLD snooping is enabled on the switch.

- Received IPv6 multicast router control packets are always flooded to the ingress VLAN, whether or not MLD snooping is enabled on the switch.
- After the discovery of the first IPv6 multicast router port, unknown IPv6 multicast data is forwarded only to the discovered router ports (before that time, all IPv6 multicast data is flooded to the ingress VLAN).

MLD Reports

The processing of MLDv1 join messages is essentially the same as with IGMPv2. When no IPv6 multicast routers are detected in a VLAN, reports are not processed or forwarded from the switch. When IPv6 multicast routers are detected and an MLDv1 report is received, an IPv6 multicast group address and an IPv6 multicast MAC address are entered in the VLAN MLD database. Then all IPv6 multicast traffic to the group within the VLAN is forwarded using this address. When MLD snooping is disabled, reports are flooded in the ingress VLAN.

When MLD snooping is enabled, MLD report suppression, called listener message suppression, is automatically enabled. With report suppression, the switch forwards the first MLDv1 report received by a group to IPv6 multicast routers; subsequent reports for the group are not sent to the routers. When MLD snooping is disabled, report suppression is disabled, and all MLDv1 reports are flooded to the ingress VLAN.

The switch also supports MLDv1 proxy reporting. When an MLDv1 MASQ is received, the switch responds with MLDv1 reports for the address on which the query arrived if the group exists in the switch on another port and if the port on which the query arrived is not the last member port for the address.

MLD Done Messages and Immediate-Leave

When the Immediate-Leave feature is enabled and a host sends an MLDv1 Done message (equivalent to an IGMP leave message), the port on which the Done message was received is immediately deleted from the group. You enable Immediate-Leave on VLANs and (as with IGMP snooping), you should only use the feature on VLANs where a single host is connected to the port. If the port was the last member of a group, the group is also deleted, and the leave information is forwarded to the detected IPv6 multicast routers.

When Immediate Leave is not enabled in a VLAN (which would be the case when there are multiple clients for a group on the same port) and a Done message is received on a port, an MASQ is generated on that port. The user can control when a port membership is removed for an existing address in terms of the number of MASQs. A port is removed from membership to an address when there are no MLDv1 reports to the address on the port for the configured number of queries.

The number of MASQs generated is configured by using the **ipv6 mld snooping last-listener-query count** global configuration command. The default number is 2.

The MASQ is sent to the IPv6 multicast address for which the Done message was sent. If there are no reports sent to the IPv6 multicast address specified in the MASQ during the switch maximum response time, the port on which the MASQ was sent is deleted from the IPv6 multicast address database. The maximum response time is the time configured by using the **ipv6 mld snooping last-listener-query-interval** global configuration command. If the deleted port is the last member of the multicast address, the multicast address is also deleted, and the switch sends the address leave information to all detected multicast routers.

Topology Change Notification Processing

When topology change notification (TCN) solicitation is enabled by using the **ipv6 mld snooping tcn query solicit** global configuration command, MLDv1 snooping sets the VLAN to flood all IPv6 multicast traffic with a configured number of MLDv1 queries before it begins sending multicast data only to selected ports. You set this value by using the **ipv6 mld snooping tcn flood query count** global configuration command. The default is to send two queries. The switch also generates MLDv1 global Done messages with valid link-local IPv6 source addresses when the switch becomes the STP root in the VLAN or when it is configured by the user. This is same as done in IGMP snooping.

MLD Snooping in Switch Stacks

The MLD IPv6 group and MAC address databases are maintained on all switches in the stack, regardless of which switch learns of an IPv6 multicast group. Report suppression and proxy reporting are done stack-wide. During the maximum response time, only one received report for a group is forwarded to the multicast routers, regardless of which switch the report arrives on.

The election of a new stack master does not affect the learning or bridging of IPv6 multicast data; bridging of IPv6 multicast data does not stop during a stack master re-election. When a new switch is added to the stack, it synchronizes the learned IPv6 multicast information from the stack master. Until the synchronization is complete, data ingressing on the newly added switch is treated as unknown multicast data.

Configuring IPv6 MLD Snooping

These sections describe how to configure IPv6 MLD snooping:

- [Default MLD Snooping Configuration, page 1-6](#)
- [MLD Snooping Configuration Guidelines, page 1-6](#)
- [Enabling or Disabling MLD Snooping, page 1-7](#)
- [Configuring a Static Multicast Group, page 1-8](#)
- [Configuring a Multicast Router Port, page 1-8](#)
- [Enabling MLD Immediate Leave, page 1-9](#)
- [Configuring MLD Snooping Queries, page 1-10](#)
- [Disabling MLD Listener Message Suppression, page 1-11](#)

Default MLD Snooping Configuration

Table 1-1 shows the default MLD snooping configuration.

Table 1-1 Default MLD Snooping Configuration

Feature	Default Setting
MLD snooping (Global)	Disabled.
MLD snooping (per VLAN)	Enabled. MLD snooping must be globally enabled for VLAN MLD snooping to take place.
IPv6 Multicast addresses	None configured.
IPv6 Multicast router ports	None configured.
MLD snooping Immediate Leave	Disabled.
MLD snooping robustness variable	Global: 2; Per VLAN: 0. Note The VLAN value overrides the global setting. When the VLAN value is 0, the VLAN uses the global count.
Last listener query count	Global: 2; Per VLAN: 0. Note The VLAN value overrides the global setting. When the VLAN value is 0, the VLAN uses the global count.
Last listener query interval	Global: 1000 (1 second); VLAN: 0. Note The VLAN value overrides the global setting. When the VLAN value is 0, the VLAN uses the global interval.
TCN query solicit	Disabled.
TCN query count	2.
MLD listener suppression	Enabled.

MLD Snooping Configuration Guidelines

When configuring MLD snooping, consider these guidelines:

- You can configure MLD snooping characteristics at any time, but you must globally enable MLD snooping by using the **ipv6 mld snooping** global configuration command for the configuration to take effect.
- When the IPv6 multicast router is a Catalyst 6500 switch and you are using extended VLANs (in the range 1006 to 4094), IPv6 MLD snooping must be enabled on the extended VLAN on the Catalyst 6500 switch in order for the Catalyst 2960 or 2960-S switch to receive queries on the VLAN. For normal-range VLANs (1 to 1005), it is not necessary to enable IPv6 MLD snooping on the VLAN on the Catalyst 6500 switch.
- MLD snooping and IGMP snooping act independently of each other. You can enable both features at the same time on the switch.
- The maximum number of address entries allowed for the switch stack is 1000.

Enabling or Disabling MLD Snooping

By default, IPv6 MLD snooping is globally disabled on the switch and enabled on all VLANs. When MLD snooping is globally disabled, it is also disabled on all VLANs. When you globally enable MLD snooping, the VLAN configuration overrides the global configuration. That is, MLD snooping is enabled only on VLAN interfaces in the default state (enabled).

You can enable and disable MLD snooping on a per-VLAN basis or for a range of VLANs, but if you globally disable MLD snooping, it is disabled in all VLANs. If global snooping is enabled, you can enable or disable VLAN snooping.

Beginning in privileged EXEC mode, follow these steps to globally enable MLD snooping on the switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ipv6 mld snooping	Globally enable MLD snooping on the switch.
Step 3	end	Return to privileged EXEC mode.
Step 4	copy running-config startup-config	(Optional) Save your entries in the configuration file.
Step 5	reload	Reload the operating system.

To globally disable MLD snooping on the switch, use the **no ipv6 mld snooping** global configuration command.

Beginning in privileged EXEC mode, follow these steps to enable MLD snooping on a VLAN.



Note

When the IPv6 multicast router is a Catalyst 6500 switch and you are using extended VLANs (in the range 1006 to 4094), IPv6 MLD snooping must be enabled on the extended VLAN on the Catalyst 6500 switch in order for the Catalyst 2960 or 2960-S switch to receive queries on the VLAN. For normal-range VLANs (1 to 1005), it is not necessary to enable IPv6 MLD snooping on the VLAN on the Catalyst 6500 switch.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ipv6 mld snooping	Globally enable MLD snooping on the switch.
Step 3	ipv6 mld snooping vlan <i>vlan-id</i>	Enable MLD snooping on the VLAN. The VLAN ID range is 1 to 1001 and 1006 to 4094. Note MLD snooping must be globally enabled for VLAN snooping to be enabled.
Step 4	end	Return to privileged EXEC mode.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable MLD snooping on a VLAN interface, use the **no ipv6 mld snooping vlan *vlan-id*** global configuration command for the specified VLAN number.

Configuring a Static Multicast Group

Hosts or Layer 2 ports normally join multicast groups dynamically, but you can also statically configure an IPv6 multicast address and member ports for a VLAN.

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	ipv6 mld snooping vlan <i>vlan-id</i> static <i>ipv6_multicast_address</i> interface <i>interface-id</i>	Statically configure a multicast group with a Layer 2 port as a member of a multicast group: <ul style="list-style-type: none"> • <i>vlan-id</i> is the multicast group VLAN ID. The VLAN ID range is 1 to 1001 and 1006 to 4094. • <i>ipv6_multicast_address</i> is the 128-bit group IPv6 address. The address must be in the form specified in RFC 2373. • <i>interface-id</i> is the member port. It can be a physical interface or a port channel (1 to 48).
Step 3	end	Return to privileged EXEC mode.
Step 4	show ipv6 mld snooping address user or show ipv6 mld snooping multicast-address vlan <i>vlan-id</i> user	Verify the static member port and the IPv6 address.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a Layer 2 port from the multicast group, use the **no ipv6 mld snooping vlan** *vlan-id* **static** *mac-address* **interface** *interface-id* global configuration command. If all member ports are removed from a group, the group is deleted.

This example shows how to statically configure an IPv6 multicast group:

```
Switch# configure terminal
Switch(config)# ipv6 mld snooping vlan 2 static FF12::3 interface gigabitethernet0/1
Switch(config)# end
```

Configuring a Multicast Router Port

Although MLD snooping learns about router ports through MLD queries and PIMv6 queries, you can also use the command-line interface (CLI) to add a multicast router port to a VLAN. To add a multicast router port (add a static connection to a multicast router), use the **ipv6 mld 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 add a multicast router port to a VLAN:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ipv6 mld 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. <ul style="list-style-type: none"> • The VLAN ID range is 1 to 1001 and 1006 to 4094. • The interface can be a physical interface or a port channel. The port-channel range is 1 to 48.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ipv6 mld snooping mrouter [vlan <i>vlan-id</i>]	Verify that IPv6 MLD 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 ipv6 mld snooping vlan *vlan-id* mrouter interface *interface-id*** global configuration command.

This example shows how to add a multicast router port to VLAN 200:

```
Switch# configure terminal
Switch(config)# ipv6 mld snooping vlan 200 mrouter interface gigabitethernet0/2
Switch(config)# exit
```

Enabling MLD Immediate Leave

When you enable MLDv1 Immediate Leave, the switch immediately removes a port from a multicast group when it detects an MLD Done message on that port. You should only use the Immediate-Leave feature when there is a single receiver present on every port in the VLAN. When there are multiple clients for a multicast group on the same port, you should not enable Immediate-Leave in a VLAN.

Beginning in privileged EXEC mode, follow these steps to enable MLDv1 Immediate Leave:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ipv6 mld snooping vlan <i>vlan-id</i> immediate-leave	Enable MLD Immediate Leave on the VLAN interface.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ipv6 mld snooping vlan <i>vlan-id</i>	Verify that Immediate Leave is enabled on the VLAN interface.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable MLD Immediate Leave on a VLAN, use the **no ipv6 mld snooping vlan *vlan-id* immediate-leave** global configuration command.

This example shows how to enable MLD Immediate Leave on VLAN 130:

```
Switch# configure terminal
Switch(config)# ipv6 mld snooping vlan 130 immediate-leave
Switch(config)# exit
```

Configuring MLD Snooping Queries

When Immediate Leave is not enabled and a port receives an MLD Done message, the switch generates MASQs on the port and sends them to the IPv6 multicast address for which the Done message was sent. You can optionally configure the number of MASQs that are sent and the length of time the switch waits for a response before deleting the port from the multicast group.

Beginning in privileged EXEC mode, follow these steps to configure MLD snooping query characteristics for the switch or for a VLAN:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ipv6 mld snooping robustness-variable <i>value</i>	(Optional) Set the number of queries that are sent before switch will deletes a listener (port) that does not respond to a general query. The range is 1 to 3; the default is 2.
Step 3	ipv6 mld snooping vlan <i>vlan-id</i> robustness-variable <i>value</i>	(Optional) Set the robustness variable on a VLAN basis, which determines the number of general queries that MLD snooping sends before aging out a multicast address when there is no MLD report response. The range is 1 to 3; the default is 0. When set to 0, the number used is the global robustness variable value.
Step 4	ipv6 mld snooping last-listener-query-count <i>count</i>	(Optional) Set the number of MASQs that the switch sends before aging out an MLD client. The range is 1 to 7; the default is 2. The queries are sent 1 second apart.
Step 5	ipv6 mld snooping vlan <i>vlan-id</i> last-listener-query-count <i>count</i>	(Optional) Set the last-listener query count on a VLAN basis. This value overrides the value configured globally. The range is 1 to 7; the default is 0. When set to 0, the global count value is used. Queries are sent 1 second apart.
Step 6	ipv6 mld snooping last-listener-query-interval <i>interval</i>	(Optional) Set the maximum response time that the switch waits after sending out a MASQ before deleting a port from the multicast group. The range is 100 to 32,768 thousands of a second. The default is 1000 (1 second).
Step 7	ipv6 mld snooping vlan <i>vlan-id</i> last-listener-query-interval <i>interval</i>	(Optional) Set the last-listener query interval on a VLAN basis. This value overrides the value configured globally. The range is 0 to 32,768 thousands of a second. The default is 0. When set to 0, the global last-listener query interval is used.
Step 8	ipv6 mld snooping tcn query solicit	(Optional) Enable topology change notification (TCN) solicitation, which means that VLANs flood all IPv6 multicast traffic for the configured number of queries before sending multicast data to only those ports requesting to receive it. The default is for TCN to be disabled.
Step 9	ipv6 mld snooping tcn flood query count <i>count</i>	(Optional) When TCN is enabled, specify the number of TCN queries to be sent. The range is from 1 to 10; the default is 2.
Step 10	end	Return to privileged EXEC mode.
Step 11	show ipv6 mld snooping querier [vlan <i>vlan-id</i>]	(Optional) Verify that the MLD snooping querier information for the switch or for the VLAN.
Step 12	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to set the MLD snooping global robustness variable to 3:

```
Switch# configure terminal
Switch(config)# ipv6 mld snooping robustness-variable 3
Switch(config)# exit
```

This example shows how to set the MLD snooping last-listener query count for a VLAN to 3:

```
Switch# configure terminal
Switch(config)# ipv6 mld snooping vlan 200 last-listener-query-count 3
Switch(config)# exit
```

This example shows how to set the MLD snooping last-listener query interval (maximum response time) to 2000 (2 seconds):

```
Switch# configure terminal
Switch(config)# ipv6 mld snooping last-listener-query-interval 2000
Switch(config)# exit
```

Disabling MLD Listener Message Suppression

MLD snooping listener message suppression is enabled by default. When it is enabled, the switch forwards only one MLD report per multicast router query. When message suppression is disabled, multiple MLD reports could be forwarded to the multicast routers.

Beginning in privileged EXEC mode, follow these steps to disable MLD listener message suppression:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no ipv6 mld snooping listener-message-suppression	Disable MLD message suppression.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ipv6 mld snooping	Verify that IPv6 MLD snooping report suppression is disabled.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To re-enable MLD message suppression, use the **ipv6 mld snooping listener-message-suppression** global configuration command.

Displaying MLD Snooping Information

You can display MLD 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 MLD snooping.

To display MLD snooping information, use one or more of the privileged EXEC commands in [Table 1-2](#).

Table 1-2 *Commands for Displaying MLD Snooping Information*

Command	Purpose
show ipv6 mld snooping [vlan <i>vlan-id</i>]	Display the MLD 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. The VLAN ID range is 1 to 1001 and 1006 to 4094.
show ipv6 mld snooping mrouter [vlan <i>vlan-id</i>]	Display information on dynamically learned and manually configured multicast router interfaces. When you enable MLD 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. The VLAN ID range is 1 to 1001 and 1006 to 4094.
show ipv6 mld snooping querier [vlan <i>vlan-id</i>]	Display information about the IPv6 address and incoming port for the most-recently received MLD query messages in the VLAN. (Optional) Enter vlan <i>vlan-id</i> to display information for a single VLAN. The VLAN ID range is 1 to 1001 and 1006 to 4094.
show ipv6 mld snooping address [vlan <i>vlan-id</i>] [count dynamic user]	Display all IPv6 multicast address information or specific IPv6 multicast address information for the switch or a VLAN. <ul style="list-style-type: none"> • Enter count to show the group count on the switch or in a VLAN. • Enter dynamic to display MLD snooping learned group information for the switch or for a VLAN. • Enter user to display MLD snooping user-configured group information for the switch or for a VLAN.
show ipv6 mld snooping multicast-address vlan <i>vlan-id</i> [<i>ipv6-multicast-address</i>]	Display MLD snooping for the specified VLAN and IPv6 multicast address.



CHAPTER 1

Configuring EtherChannels and Link-State Tracking



Note

To use link-state tracking, the switch must be running the LAN Base image.



Note

This chapter describes how to configure EtherChannels on the Catalyst 2960 and 2960-S switches. 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. This chapter also describes how to configure link-state tracking. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack. Stacking is supported only on Catalyst 2960-S switches running the LAN base image.



Note

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release.

- [Understanding EtherChannels, page 1-1](#)
- [Configuring EtherChannels, page 1-11](#)
- [Displaying EtherChannel, PAgP, and LACP Status, page 1-20](#)
- [Understanding Link-State Tracking, page 1-20](#)
- [Configuring Link-State Tracking, page 1-23](#)

Understanding EtherChannels


- [EtherChannel Overview, page 1-2](#)
- [Port-Channel Interfaces, page 1-4](#)
- [Port Aggregation Protocol, page 1-5](#)
- [Link Aggregation Control Protocol, page 1-7](#)

- [EtherChannel On Mode, page 1-8](#)
- [Load Balancing and Forwarding Methods, page 1-8](#)
- [EtherChannel and Switch Stacks, page 1-10](#)

EtherChannel Overview

An EtherChannel consists of individual Fast Ethernet or Gigabit Ethernet links bundled into a single logical link as shown in [Figure 1-1](#).

Figure 1-1 Typical EtherChannel Configuration



The EtherChannel provides full-duplex bandwidth up to 800 Mb/s (Fast EtherChannel) or 8 Gb/s (Gigabit EtherChannel) between your switch and another switch or host. Each EtherChannel can consist of up to eight compatibly configured Ethernet ports.

All ports in each EtherChannel must be configured as Layer 2 ports. The number of EtherChannels is limited to six.

For more information, see the [“EtherChannel Configuration Guidelines”](#) section on page 1-11.

You can configure an EtherChannel in one of these modes: Port Aggregation Protocol (PAgP), Link Aggregation Control Protocol (LACP), or On. Configure both ends of the EtherChannel in the same mode:

- When you configure one end of an EtherChannel in either PAgP or LACP mode, the system negotiates with the other end of the channel to determine which ports should become active. Incompatible ports are put into an independent state and continue to carry data traffic as would any other single link. The port configuration does not change, but the port does not participate in the EtherChannel.
- When you configure an EtherChannel in the **on** mode, no negotiations take place. The switch forces all compatible ports to become active in the EtherChannel. The other end of the channel (on the other switch) must also be configured in the **on** mode; otherwise, packet loss can occur.

You can create an EtherChannel on a standalone switch, on a single switch in the stack, or on multiple switches in the stack (known as cross-stack EtherChannel). See [Figure 1-2](#) and [Figure 1-3](#).

If a link within an EtherChannel fails, traffic previously carried over that failed link moves to the remaining links within the EtherChannel. If traps are enabled on the switch, a trap is sent for a failure that identifies 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.

Figure 1-2 Single-Switch EtherChannel

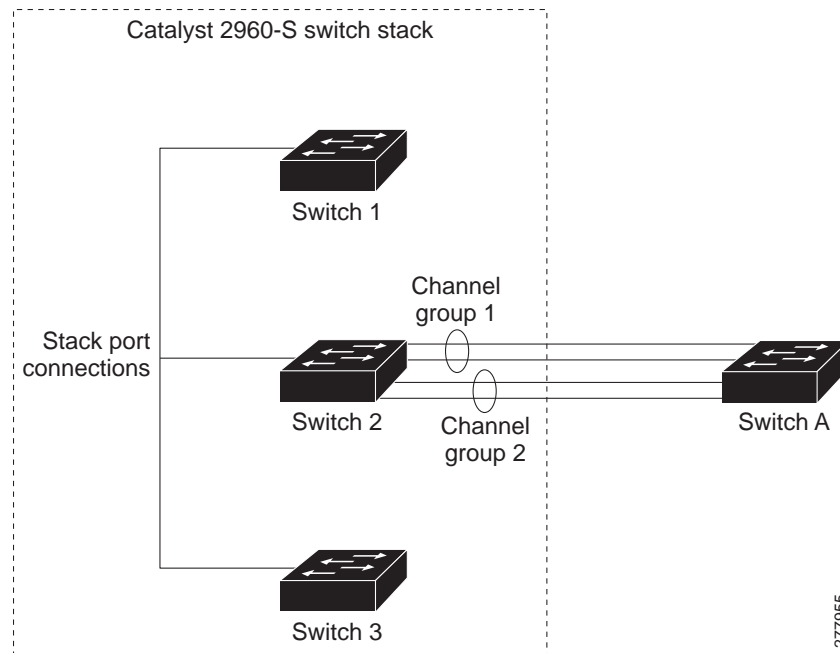



Figure 1-3 Cross-Stack EtherChannel



277964


Port-Channel Interfaces

When you create a Layer 2 EtherChannel, a port-channel logical interface is involved. You can create the EtherChannel in these ways:

- Use the **channel-group** interface configuration command. This command automatically creates the port-channel logical interface when the channel group gets its first physical port. The **channel-group** command binds the physical (10/100/1000 ports) and the logical ports together as shown in Figure 1-4.
- Use the **interface port-channel** *port-channel-number* global configuration command to manually create the port-channel logical interface. Then use the **channel-group** *channel-group-number* interface configuration 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.

Each EtherChannel has a port-channel logical interface numbered from 1 to 6. This port-channel interface number corresponds to the one specified with the **channel-group** interface configuration command.

Figure 1-4 Relationship of Physical Ports, Logical Port Channels, and Channel Groups



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. 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.

You can use PAgP only in single-switch EtherChannel configurations; PAgP cannot be enabled on cross-stack EtherChannels. PAgP dynamically groups similarly configured ports on a single switch in the stack into a single logical link. For more information, see the [“EtherChannel Configuration Guidelines”](#) section on page 1-11.

PAgP Modes

Table 1-1 shows the user-configurable EtherChannel PAgP modes for the **channel-group** interface configuration command.

Table 1-1 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. This mode is not supported when the EtherChannel members are from different switches in the switch stack (cross-stack EtherChannel).
desirable	Places a port into an active negotiating state, in which the port starts negotiations with other ports by sending PAgP packets. This mode is not supported when the EtherChannel members are from different switches in the switch stack (cross-stack EtherChannel).

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.

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 Virtual Switches and Dual-Active Detection

A virtual switch can be two or more Catalyst 6500 core switches connected by virtual switch links (VSLs) that carry control and data traffic between them. One of the switches is in active mode. The others are in standby mode. For redundancy, remote switches, such as Catalyst 2960 or 2960-S switches, are connected to the virtual switch by remote satellite links (RSLs).



Note

Only a Catalyst 2960 switch running the LAN Base image can be remote switch.

If the VSL between two switches fails, one switch does not know the status of the other. Both switches could change to the active mode, causing a *dual-active situation* in the network with duplicate configurations (including duplicate IP addresses and bridge identifiers). The network might go down.

To prevent a dual-active situation, the core switches send PAgP protocol data units (PDUs) through the RSLs to the remote switches. The PAgP PDUs identify the active switch, and the remote switches forward the PDUs to core switches so that the core switches are in sync. If the active switch fails or resets, the standby switch takes over as the active switch. If the VSL goes down, one core switch knows the status of the other and does not change state.

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 IEEE 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 1-2 shows the user-configurable EtherChannel LACP modes for the **channel-group** interface configuration command.

Table 1-2 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.

Both the **active** and **passive LACP** modes enable ports to negotiate with partner ports to 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.

EtherChannel On Mode

EtherChannel **on** mode can be used to manually configure an EtherChannel. The **on** mode forces a port to join an EtherChannel without negotiations. The **on** mode can be useful if the remote device does not support PAgP or LACP. In the **on** mode, a usable EtherChannel exists only when the switches at both ends of the link are configured in the **on** mode.

Ports that are configured in the **on** mode in the same channel group must have compatible port characteristics, such as speed and duplex. Ports that are not compatible are suspended, even though they are configured in the **on** mode.



Caution

You should use care when using the **on** mode. This is a manual configuration, and ports on both ends of the EtherChannel must have the same configuration. If the group is misconfigured, packet loss or spanning-tree loops can occur.

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, packets are sent to an EtherChannel and distributed across the EtherChannel ports, 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 1-5](#), an EtherChannel from a switch that is aggregating data from 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 only going 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 1-5 Load Distribution and Forwarding Methods



EtherChannel and Switch Stacks

If a stack member that has ports participating in an EtherChannel fails or leaves the stack, the stack master removes the failed stack member switch ports from the EtherChannel. The remaining ports of the EtherChannel, if any, continue to provide connectivity.

When a switch is added to an existing stack, the new switch receives the running configuration from the stack master and updates itself with the EtherChannel-related stack configuration. The stack member also receives the operational information (the list of ports that are up and are members of a channel).

When two stacks merge that have EtherChannels configured between them, self-looped ports result. Spanning tree detects this condition and acts accordingly. Any PAgP or LACP configuration on a winning switch stack is not affected, but the PAgP or LACP configuration on the losing switch stack is lost after the stack reboots.

With PAgP, if the stack master fails or leaves the stack, a new stack master is elected. A spanning-tree reconvergence is not triggered unless there is a change in the EtherChannel bandwidth. The new stack master synchronizes the configuration of the stack members to that of the stack master. The PAgP configuration is not affected after a stack master change unless the EtherChannel has ports residing on the old stack master.

With LACP, the system-id uses the stack MAC address from the stack master, and if the stack master changes, the LACP system-id can change. If the LACP system-id changes, the entire EtherChannel will flap, and there will be an STP reconvergence. Use the **stack-mac persistent timer** command to control whether or not the stack MAC address changes during a master failover.

For more information about switch stacks, see [Chapter 1, “Managing Switch Stacks.”](#)

Configuring EtherChannels

- [Default EtherChannel Configuration, page 1-11](#)
- [EtherChannel Configuration Guidelines, page 1-11](#)
- [Configuring Layer 2 EtherChannels, page 1-13](#) (required)
- [Configuring EtherChannel Load Balancing, page 1-15](#) (optional)
- [Configuring the PAgP Learn Method and Priority, page 1-16](#) (optional)
- [Configuring LACP Hot-Standby Ports, page 1-18](#) (optional)



Note

Make sure that the ports are correctly configured. For more information, see the “[EtherChannel Configuration Guidelines](#)” section on [page 1-11](#).



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 1-3 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 or switch stack MAC address.
Load balancing	Load distribution on the switch is based on the source-MAC address of the incoming packet.

EtherChannel Configuration Guidelines



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

If improperly configured, some EtherChannel ports are automatically disabled to avoid network loops and other problems. Follow these guidelines to avoid configuration problems:

- Do not try to configure more than 6 EtherChannels on the switch stack.
- 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 or on different switches in the stack. Individual EtherChannel groups can run either PAgP or LACP, but they cannot interoperate.
- Do not configure a Switched Port Analyzer (SPAN) destination port 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 or a not-yet-active member of an EtherChannel as an IEEE 802.1x port. If you try to enable IEEE 802.1x on an EtherChannel port, an error message appears, and IEEE 802.1x is not enabled.
- If EtherChannels are configured on switch interfaces, remove the EtherChannel configuration from the interfaces before globally enabling IEEE 802.1x on a switch by using the **dot1x system-auth-control** global configuration command.
- Do not enable link-state tracking on individual interfaces that will be part of a downstream Etherchannel interface.
- 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 IEEE 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 cross-stack EtherChannel configurations, ensure that all ports targeted for the EtherChannel are either configured for LACP or are manually configured to be in the channel group using the **channel-group** *channel-group-number* **mode on** interface configuration command. The PAgP protocol is not supported on cross-stack EtherChannels.
- If cross-stack EtherChannel is configured and the switch stack partitions, loops and forwarding misbehaviors can occur.

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.

If you enabled PAgP on a port in the **auto** or **desirable** mode, you must reconfigure it for either the **on** mode or the LACP mode before adding this port to a cross-stack EtherChannel. PAgP does not support cross-stack EtherChannels.

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, 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 }	<p>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 6.</p> <p>For mode, select one of these keywords:</p> <ul style="list-style-type: none"> • 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. The auto keyword is not supported when EtherChannel members are from different switches in the switch stack. • 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. The desirable keyword is not supported when EtherChannel members are from different switches in the switch stack. • on—Forces the port to channel without PAgP or LACP. In the on mode, an 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. <p>For information on compatible modes for the switch and its partner, see the “PAgP Modes” section on page 1-6 and the “LACP Modes” section on page 1-7.</p>
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 on a switch. 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 gigabitethernet0/1 -2
Switch(config-if-range)# switchport mode access
Switch(config-if-range)# switchport access vlan 10
Switch(config-if-range)# channel-group 5 mode desirable non-silent
Switch(config-if-range)# end
```

This example shows how to configure an EtherChannel on a switch. 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 gigabitethernet0/1 -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
```

This example shows how to configure a cross-stack EtherChannel. It uses LACP passive mode and assigns two ports on stack member 2 and one port on stack member 3 as static-access ports in VLAN 10 to channel 5:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet2/0/4 -5
Switch(config-if-range)# switchport mode access
Switch(config-if-range)# switchport access vlan 10
Switch(config-if-range)# channel-group 5 mode active
Switch(config-if-range)# exit
Switch(config)# interface gigabitethernet3/0/3
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan 10
Switch(config-if)# channel-group 5 mode active
Switch(config-if)# exit
```

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 1-8](#).

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: <ul style="list-style-type: none"> • 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.

When the link partner of the switch is a physical learner (such as a Catalyst 1900 series switch), we recommend that you configure the Catalyst 2960 or 2960-S 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. Only use the **pagp learn-method** command 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.
Step 3	pagp learn-method <i>physical-port</i>	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 1-15. The learning method must be configured the same at both ends of the link.
Step 4	pagp port-priority <i>priority</i>	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.

	Command	Purpose
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. To every link between systems that operate LACP, the software assigns a unique priority made up of these elements (in priority order):

- LACP system priority
- System ID (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.

Determining which ports are active and which are hot standby is a two-step procedure. First the system with a numerically lower system priority and system-id is placed in charge of the decision. Next, that system decides which ports are active and which are hot standby, based on its values for port priority and port number. The port-priority and port-number values for the other system are not used.

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 1-18](#) and the [“Configuring the LACP Port Priority” section on page 1-19](#).

Configuring the LACP System Priority

You can configure the system priority for all 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 default 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.

	Command	Purpose
Step 5	show running-config or show lacp [channel-group-number] 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

Table 1-4 Commands for Displaying EtherChannel, PAgP, and LACP Status

Command	Description
show etherchannel [channel-group-number { 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 [channel-group-number] { counters internal neighbor }	Displays PAgP information such as traffic information, the internal PAgP configuration, and neighbor information.
show pagp [channel-group-number] dual-active	Displays the dual-active detection status.
show lacp [channel-group-number] { 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, see the command reference for this release.

Understanding Link-State Tracking



Note

To use Link-state tracking, the switch must be running the LAN Base image.

Link-state tracking, also known as trunk failover, is a feature that binds the link state of multiple interfaces. For example, link-state tracking provides redundancy in the network when used with server NIC adapter teaming. When the server network adapters are configured in a primary or secondary relationship known as teaming, if the link is lost on the primary interface, connectivity is transparently changed to the secondary interface.

**Note**

An interface can be an aggregation of ports (an EtherChannel), or a single physical port in access or trunk mode.

Figure 1-6 on page 1-23 shows a network configured with link-state tracking. To enable link-state tracking, create a *link-state group*, and specify the interfaces that are assigned to the link-state group. In a link-state group, these interfaces are bundled together. The *downstream interfaces* are bound to the *upstream interfaces*. Interfaces connected to servers are referred to as downstream interfaces, and interfaces connected to distribution switches and network devices are referred to as upstream interfaces.

The configuration in Figure 1-6 ensures that the network traffic flow is balanced as follows:

- For links to switches and other network devices
 - Server 1 and server 2 use switch A for primary links and switch B for secondary links.
 - Server 3 and server 4 use switch B for primary links and switch A for secondary links.
- Link-state group 1 on switch A
 - Switch A provides primary links to server 1 and server 2 through link-state group 1. Port 1 is connected to server 1, and port 2 is connected to server 2. Port 1 and port 2 are the downstream interfaces in link-state group 1.
 - Port 5 and port 6 are connected to distribution switch 1 through link-state group 1. Port 5 and port 6 are the upstream interfaces in link-state group 1.
- Link-state group 2 on switch A
 - Switch A provides secondary links to server 3 and server 4 through link-state group 2. Port 3 is connected to server 3, and port 4 is connected to server 4. Port 3 and port 4 are the downstream interfaces in link-state group 2.
 - Port 7 and port 8 are connected to distribution switch 2 through link-state group 2. Port 7 and port 8 are the upstream interfaces in link-state group 2.
- Link-state group 2 on switch B
 - Switch B provides primary links to server 3 and server 4 through link-state group 2. Port 3 is connected to server 3, and port 4 is connected to server 4. Port 3 and port 4 are the downstream interfaces in link-state group 2.
 - Port 5 and port 6 are connected to distribution switch 2 through link-state group 2. Port 5 and port 6 are the upstream interfaces in link-state group 2.
- Link-state group 1 on switch B
 - Switch B provides secondary links to server 1 and server 2 through link-state group 1. Port 1 is connected to server 1, and port 2 is connected to server 2. Port 1 and port 2 are the downstream interfaces in link-state group 1.
 - Port 7 and port 8 are connected to distribution switch 1 through link-state group 1. Port 7 and port 8 are the upstream interfaces in link-state group 1.

In a link-state group, the upstream ports can become unavailable or lose connectivity because the distribution switch or router fails, the cables are disconnected, or the link is lost. These are the interactions between the downstream and upstream interfaces when link-state tracking is enabled:

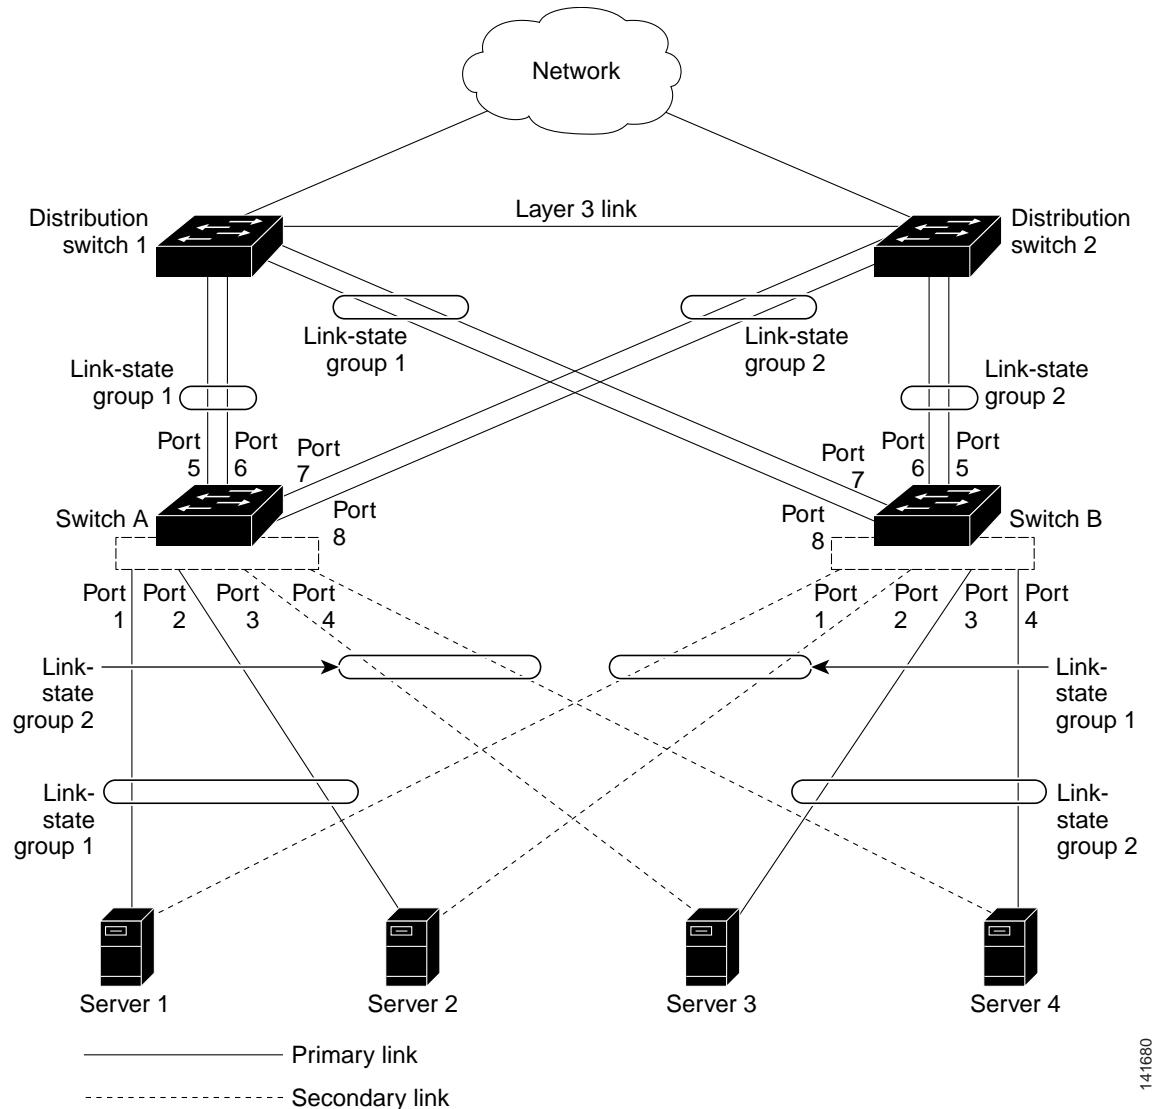
- If any of the upstream interfaces are in the link-up state, the downstream interfaces can change to or remain in the link-up state.
- If all of the upstream interfaces become unavailable, link-state tracking automatically puts the downstream interfaces in the error-disabled state. Connectivity to and from the servers is automatically changed from the primary server interface to the secondary server interface.

As an example of a connectivity change from link-state group 1 to link-state group 2 on switch A, see [Figure 1-6 on page 1-23](#). If the upstream link for port 6 is lost, the link states of downstream ports 1 and 2 do not change. However, if the link for upstream port 5 is also lost, the link state of the downstream ports changes to the link-down state. Connectivity to server 1 and server 2 is then changed from link-state group 1 to link-state group 2. The downstream ports 3 and 4 do not change state because they are in link-group 2.

- If the link-state group is configured, link-state tracking is disabled, and the upstream interfaces lose connectivity, the link states of the downstream interfaces remain unchanged. The server does not recognize that upstream connectivity has been lost and does not failover to the secondary interface.

You can recover a downstream interface link-down condition by removing the failed downstream port from the link-state group. To recover multiple downstream interfaces, disable the link-state group.

Figure 1-6 Typical Link-State Tracking Configuration



141680

Configuring Link-State Tracking

- [Default Link-State Tracking Configuration, page 1-23](#)
- [Link-State Tracking Configuration Guidelines, page 1-24](#)
- [Configuring Link-State Tracking, page 1-24](#)
- [Displaying Link-State Tracking Status, page 1-25](#)

Default Link-State Tracking Configuration

There are no link-state groups defined, and link-state tracking is not enabled for any group.

Link-State Tracking Configuration Guidelines

Follow these guidelines to avoid configuration problems:

- An interface that is defined as an upstream interface cannot also be defined as a downstream interface in the same or a different link-state group. The reverse is also true.
- Do not enable link-state tracking on individual interfaces that will be part of a downstream Etherchannel interface.
- An interface cannot be a member of more than one link-state group.
- You can configure only two link-state groups per switch.

Configuring Link-State Tracking

Beginning in privileged EXEC mode, follow these steps to configure a link-state group and to assign an interface to a group:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	link state track <i>number</i>	Create a link-state group, and enable link-state tracking. The group number can be 1 to 2; the default is 1.
Step 3	interface <i>interface-id</i>	Specify a physical interface or range of interfaces to configure, and enter interface configuration mode. Valid interfaces include switch ports in access or trunk mode (IEEE 802.1q), routed ports, or multiple ports bundled into an upstream EtherChannel interface (static, PAgP, or LACP), also in trunk mode. Note Do not enable link-state tracking on individual interfaces that will be part of a downstream Etherchannel interface.
Step 4	link state group [<i>number</i>] { upstream downstream }	Specify a link-state group, and configure the interface as either an upstream or downstream interface in the group. The group number can be 1 to 2; the default is 1.
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.

This example shows how to create a link-state group and configure the interfaces:

```
Switch# configure terminal
Switch(config)# link state track 1
Switch(config)# interface range gigabitethernet0/21 -22
Switch(config-if)# link state group 1 upstream
Switch(config-if)# interface gigabitethernet0/1
Switch(config-if)# link state group 1 downstream
Switch(config-if)# interface gigabitethernet0/3
Switch(config-if)# link state group 1 downstream
Switch(config-if)# interface gigabitethernet0/5
Switch(config-if)# link state group 1 downstream
Switch(config-if)# end
```

To disable a link-state group, use the **no link state track *number*** global configuration command.

Displaying Link-State Tracking Status

Use the **show link state group** command to display the link-state group information. Enter this command without keywords to display information about all link-state groups. Enter the group number to display information specific to the group. Enter the detail keyword to display detailed information about the group.

This is an example of output from the **show link state group 1** command:

```
Switch> show link state group 1
```

```
Link State Group: 1      Status: Enabled, Down
```

This is an example of output from the **show link state group detail** command:

```
Switch> show link state group detail
```

```
(Up):Interface up      (Dwn):Interface Down  (Dis):Interface disabled
```

```
Link State Group: 1 Status: Enabled, Down
```

```
Upstream Interfaces : Gi0/15(Dwn) Gi0/16(Dwn)
```

```
Downstream Interfaces : Gi0/11(Dis) Gi0/12(Dis) Gi0/13(Dis) Gi0/14(Dis)
```

```
Link State Group: 2 Status: Enabled, Down
```

```
Upstream Interfaces : Gi1/0/15(Dwn) Gi1/0/16(Dwn) Gi1/0/17(Dwn)
```

```
Downstream Interfaces : Gi1/0/11(Dis) Gi1/0/12(Dis) Gi1/0/13(Dis) Gi1/0/14(Dis)
```

```
(Up):Interface up      (Dwn):Interface Down  (Dis):Interface disabled
```

For detailed information about the fields in the display, see the command reference for this release.



CHAPTER 1

Troubleshooting

This chapter describes how to identify and resolve software problems related to the Cisco IOS software on the Catalyst 2960 and 2960-S switches. Depending on the nature of the problem, you can use the command-line interface (CLI), the device manager, or Network Assistant to identify and solve problems. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

Additional troubleshooting information, such as LED descriptions, is provided in the hardware installation guide.



Note For complete syntax and usage information for the commands used in this chapter, see the command reference for this release and the *Cisco IOS Commands Master List, Release 12.4* on Cisco.com.

This chapter consists of these sections:

- [Recovering from a Software Failure, page 1-2](#)
- [Recovering from a Lost or Forgotten Password, page 1-3](#)
- [Preventing Switch Stack Problems, page 1-7](#)
- [Recovering from a Command Switch Failure, page 1-8](#)
- [Recovering from Lost Cluster Member Connectivity, page 1-12](#)



Note Recovery procedures require that you have physical access to the switch.

- [Preventing Autonegotiation Mismatches, page 1-12](#)
- [Troubleshooting Power over Ethernet Switch Ports, page 1-12](#)
- [SFP Module Security and Identification, page 1-13](#)
- [Monitoring SFP Module Status, page 1-14](#)
- [Using Ping, page 1-14](#)
- [Using Layer 2 Traceroute, page 1-15](#)
- [Using IP Traceroute, page 1-17](#)
- [Using TDR, page 1-19](#)
- [Using Debug Commands, page 1-19](#)

- Using the show platform forward Command, page 1-21
- Using the crashinfo Files, page 1-23
- Using On-Board Failure Logging, page 1-24
- Memory Consistency Check Routines, page 1-26
- Troubleshooting Tables, page 1-27

Recovering from a Software Failure

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 that 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, see the release notes.

Step 2 Extract the bin file from the tar file.

- If you are using Windows, use a zip program that can read 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.

```
unix-1% tar -tvf image_filename.tar
```

2. Locate the bin file, and extract it by using the `tar -xvf <image_filename.tar> <image_filename.bin>` UNIX command.

```
unix-1% tar -xvf image_filename.tar image_filename.bin
x c2960-lanbase-mz.122-25.FX/c2960-lanbase-mz.122-25.FX.bin, 2928176 bytes, 5720
tape blocks
```

3. Verify that the bin file was extracted by using the `ls -l <image_filename.bin>` UNIX command.

```
unix-1% ls -l image_filename.bin
-rw-r--r--  1 boba      2928176 Apr 21 12:01
c2960-lanbase-mz.122-25.FX/c2960-lanbase-mz.122-25.FX.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 or to the switch stack.

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.

These sections describes how to recover a forgotten or lost switch password:

- [Procedure with Password Recovery Enabled, page 1-4](#)
- [Procedure with Password Recovery Disabled, page 1-6](#)

You enable or disable password recovery by using the **service password-recovery** global configuration command. When you enter the **service password-recovery** or **no service password-recovery** command on the stack master, it is propagated throughout the stack and applied to all switches in the stack.



Note Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

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 standalone switch or the entire switch stack.
- Step 4** Reconnect the power cord to the standalone switch or the stack master and, within 15 seconds, press the **Mode** button while the System LED is still flashing green. Continue pressing the **Mode** button until the System LED turns briefly amber and then solid green; then release the **Mode** button.

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
```

go to the [“Procedure with Password Recovery Enabled”](#) section on page 1-4, and follow the steps.

- If you see a message that begins with this:

```
The password-recovery mechanism has been triggered, but is currently disabled.
```

go to the [“Procedure with Password Recovery Disabled”](#) section on page 1-6, and follow the steps.

- Step 5** After recovering the password, reload the standalone switch or the stack master:

```
Switch> reload
 slot <stack-master-member-number>
Proceed with reload? [confirm] y
```

- Step 6** Power on the rest of the switch stack.
-

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:

```
Directory of flash:
 13 drwx          192  Mar 01 1993 22:30:48 c3560c2960-lanbase-mz.122-25.FX
 11 -rwx          5825  Mar 01 1993 22:31:59 config.text
 18 -rwx           720  Mar 01 1993 02:21:30 vlan.dat
```

```
16128000 bytes total (10003456 bytes free)
```

- 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 up 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 stack:

```
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:

```
Directory of flash:
```

```
13 drwx          192  Mar 01 1993 22:30:48 c2960-lanbase-mz.122-25.FX.0
```

```
16128000 bytes total (10003456 bytes free)
```

Step 4 Boot up 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#
```



Note Before continuing to Step 9, power on any connected stack members and wait until they have completely initialized.
Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

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 Switch Stack Problems



Note

- Make sure that the switches that you add to or remove from the switch stack are powered off. For all powering considerations in switch stacks, see the “Switch Installation” chapter in the hardware installation guide.

- After adding or removing stack members, make sure that the switch stack is operating at full bandwidth (32 Gb/s). Press the Mode button on a stack member until the Stack mode LED is on. The last two port LEDs on the switch should be green. Depending on the switch model, the last two ports are either 10/100/1000 ports or small form-factor pluggable (SFP) module ports. If one or both of the last two port LEDs are not green, the stack is not operating at full bandwidth.
- We recommend using only one CLI session when managing the switch stack. Be careful when using multiple CLI sessions to the stack master. Commands that you enter in one session are not displayed in the other sessions. Therefore, it is possible that you might not be able to identify the session from which you entered a command.
- Manually assigning stack member numbers according to the placement of the switches in the stack can make it easier to remotely troubleshoot the switch stack. However, you need to remember that the switches have manually assigned numbers if you add, remove, or rearrange switches later. Use the **switch** *current-stack-member-number* **renumber** *new-stack-member-number* global configuration command to manually assign a stack member number. For more information about stack member numbers, see the “[Member Numbers](#)” section on page 1-6.

If you replace a stack member with an identical model, the new switch functions with the exact same configuration as the replaced switch. This is also assuming the new switch is using the same member number as the replaced switch.

Removing powered-on stack members causes the switch stack to divide (partition) into two or more switch stacks, each with the same configuration. If you want the switch stacks to remain separate, change the IP address or addresses of the newly created switch stacks. To recover from a partitioned switch stack:

1. Power off the newly created switch stacks.
2. Reconnect them to the original switch stack through their StackWise ports.
3. Power on the switches.

For the commands that you can use to monitor the switch stack and its members, see the “[Displaying Stack Information](#)” section on page 1-22.

Recovering from a Command Switch Failure

This section describes how to recover from a failed command switch. You can configure a redundant command switch group by using the Hot Standby Router Protocol (HSRP). For more information, see [Chapter 1, “Clustering Switches.”](#) Also see the *Getting Started with Cisco Network Assistant*, available on Cisco.com.



Note

HSRP is the preferred method for supplying redundancy to a cluster.

If you have not configured a standby command switch, and your command switch loses power or fails in some other way, management contact with the member switches is lost, and you must install a new command switch. However, connectivity between switches that are still connected is not affected, and the member switches forward packets as usual. You can manage the members as standalone switches through the console port, or, if they have IP addresses, through the other management interfaces.

You can prepare for a command switch failure by assigning an IP address to a member switch or another switch that is command-capable, making a note of the command-switch password, and cabling your cluster to provide redundant connectivity between the member switches and the replacement command switch. These sections describe two solutions for replacing a failed command switch:

- [Replacing a Failed Command Switch with a Cluster Member, page 1-9](#)
- [Replacing a Failed Command Switch with Another Switch, page 1-10](#)

These recovery procedures require that you have physical access to the switch.

For information on command-capable switches, see the release notes.

Replacing a Failed Command Switch with a Cluster Member

To replace a failed command switch with a command-capable member in the same cluster, follow these steps:

- Step 1** Disconnect the command switch from the member switches, and physically remove it from the cluster.
- Step 2** Insert the member switch in place of the failed command switch, and duplicate its connections to the cluster members.
- Step 3** Start a CLI session on the new command switch.

You can access the CLI by using the console port or, if an IP address has been assigned to the switch, by using Telnet. For details about using the console port, see the switch hardware installation guide.

- Step 4** At the switch prompt, enter privileged EXEC mode:

```
Switch> enable
Switch#
```

- Step 5** Enter the password of the *failed command switch*.

- Step 6** Enter global configuration mode.

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
```

- Step 7** Remove the member switch from the cluster.

```
Switch(config)# no cluster commander-address
```

- Step 8** Return to privileged EXEC mode.

```
Switch(config)# end
Switch#
```

- Step 9** Use the setup program to configure the switch IP information. This program prompts you for IP address information and passwords. From privileged EXEC mode, enter **setup**, and press **Return**.

```
Switch# setup
--- System Configuration Dialog ---
Continue with configuration dialog? [yes/no]: y
```

At any point you may enter a question mark '?' for help.
Use ctrl-c to abort configuration dialog at any prompt.
Default settings are in square brackets '[]'.

Basic management setup configures only enough connectivity for management of the system, extended setup will ask you to configure each interface on the system

Would you like to enter basic management setup? [yes/no]:

Step 10 Enter **Y** at the first prompt.

The prompts in the setup program vary depending on the member switch that you selected to be the command switch:

```
Continue with configuration dialog? [yes/no]: y
or
```

```
Configuring global parameters:
```

If this prompt does not appear, enter **enable**, and press **Return**. Enter **setup**, and press **Return** to start the setup program.

Step 11 Respond to the questions in the setup program.

When prompted for the hostname, recall that on a command switch, the hostname is limited to 28 characters; on a member switch to 31 characters. Do not use *-n*, where *n* is a number, as the last characters in a hostname for any switch.

When prompted for the Telnet (virtual terminal) password, recall that it can be from 1 to 25 alphanumeric characters, is case sensitive, allows spaces, but ignores leading spaces.

Step 12 When prompted for the **enable secret** and **enable** passwords, enter the passwords of the *failed command switch* again.

Step 13 When prompted, make sure to enable the switch as the cluster command switch, and press **Return**.

Step 14 When prompted, assign a name to the cluster, and press **Return**.

The cluster name can be 1 to 31 alphanumeric characters, dashes, or underscores.

Step 15 After the initial configuration displays, verify that the addresses are correct.

Step 16 If the displayed information is correct, enter **Y**, and press **Return**.

If this information is not correct, enter **N**, press **Return**, and begin again at Step 9.

Step 17 Start your browser, and enter the IP address of the new command switch.

Step 18 From the Cluster menu, select **Add to Cluster** to display a list of candidate switches to add to the cluster.

Replacing a Failed Command Switch with Another Switch

To replace a failed command switch with a switch that is command-capable but not part of the cluster, follow these steps:

Step 1 Insert the new switch in place of the failed command switch, and duplicate its connections to the cluster members.

Step 2 Start a CLI session on the new command switch.

You can access the CLI by using the console port or, if an IP address has been assigned to the switch, by using Telnet. For details about using the console port, see the switch hardware installation guide.

Step 3 At the switch prompt, enter privileged EXEC mode:

```
Switch> enable
Switch#
```

Step 4 Enter the password of the *failed command switch*.

Step 5 Use the setup program to configure the switch IP information.

This program prompts you for IP address information and passwords. From privileged EXEC mode, enter **setup**, and press **Return**.

```
Switch# setup
--- System Configuration Dialog ---
Continue with configuration dialog? [yes/no]: y

At any point you may enter a question mark '?' for help.
Use ctrl-c to abort configuration dialog at any prompt.
Default settings are in square brackets '[]'.

Basic management setup configures only enough connectivity
for management of the system, extended setup will ask you
to configure each interface on the system

Would you like to enter basic management setup? [yes/no]:
```

Step 6 Enter **Y** at the first prompt.

The prompts in the setup program vary depending on the switch you selected to be the command switch:

```
Continue with configuration dialog? [yes/no]: y
```

or

```
Configuring global parameters:
```

If this prompt does not appear, enter **enable**, and press **Return**. Enter **setup**, and press **Return** to start the setup program.

Step 7 Respond to the questions in the setup program.

When prompted for the hostname, recall that on a command switch, the hostname is limited to 28 characters. Do not use *-n*, where *n* is a number, as the last character in a hostname for any switch.

When prompted for the Telnet (virtual terminal) password, recall that it can be from 1 to 25 alphanumeric characters, is case sensitive, allows spaces, but ignores leading spaces.

Step 8 When prompted for the **enable secret** and **enable** passwords, enter the passwords of the *failed command switch* again.

Step 9 When prompted, make sure to enable the switch as the cluster command switch, and press **Return**.

Step 10 When prompted, assign a name to the cluster, and press **Return**.

The cluster name can be 1 to 31 alphanumeric characters, dashes, or underscores.

Step 11 When the initial configuration displays, verify that the addresses are correct.

Step 12 If the displayed information is correct, enter **Y**, and press **Return**.

If this information is not correct, enter **N**, press **Return**, and begin again at Step 9.

Step 13 Start your browser, and enter the IP address of the new command switch.

Step 14 From the Cluster menu, select **Add to Cluster** to display a list of candidate switches to add to the cluster.

Recovering from Lost Cluster Member Connectivity

Some configurations can prevent the command switch from maintaining contact with member switches. If you are unable to maintain management contact with a member, and the member switch is forwarding packets normally, check for these conflicts:

- A member switch (Catalyst 3750, Catalyst 3560, Catalyst 3550, Catalyst 3500 XL, Catalyst 2970, Catalyst 2960, Catalyst 2950, Catalyst 2900 XL, Catalyst 2820, and Catalyst 1900 switch) cannot connect to the command switch through a port that is defined as a network port.
- Catalyst 3500 XL, Catalyst 2900 XL, Catalyst 2820, and Catalyst 1900 member switches must connect to the command switch through a port that belongs to the same management VLAN.
- A member switch (Catalyst 3750, Catalyst 3560, Catalyst 3550, Catalyst 2970, Catalyst 2960, Catalyst 2950, Catalyst 3500 XL, Catalyst 2900 XL, Catalyst 2820, and Catalyst 1900 switch) connected to the command switch through a secured port can lose connectivity if the port is disabled because of a security violation.

Preventing Autonegotiation Mismatches

The IEEE 802.3ab autonegotiation protocol manages the switch settings for speed (10 Mb/s, 100 Mb/s, and 1000 Mb/s, excluding SFP module ports) and duplex (half or full). There are situations when this protocol can incorrectly align these settings, reducing performance. A mismatch occurs under these circumstances:

- A manually set speed or duplex parameter is different from the manually set speed or duplex parameter on the connected port.
- A port is set to autonegotiate, and the connected port is set to full duplex with no autonegotiation.

To maximize switch performance and ensure a link, follow one of these guidelines when changing the settings for duplex and speed:

- Let both ports autonegotiate both speed and duplex.
- Manually set the speed and duplex parameters for the ports on both ends of the connection.



Note

If a remote device does not autonegotiate, configure the duplex settings on the two ports to match. The speed parameter can adjust itself even if the connected port does not autonegotiate.

Troubleshooting Power over Ethernet Switch Ports

These sections describe how to troubleshoot Power over Ethernet (PoE) ports.

Disabled Port Caused by Power Loss

If a powered device (such as a Cisco IP Phone 7910) that is connected to a PoE switch port and is powered by an AC power source loses power from the AC power source, the device might enter an error-disabled state. To recover from an error-disabled state, enter the **shutdown** interface configuration command, and then enter the **no shutdown** interface command. You can also configure automatic

recovery on the switch to recover from the error-disabled state. The **errdisable recovery cause loopback** and the **errdisable recovery interval** *seconds* global configuration commands automatically take the interface out of the error-disabled state after the specified period of time.

Use these commands, described in the command reference for this release, to monitor the PoE port status:

- **show controllers power inline** privileged EXEC command
- **show power inline** privileged EXEC command
- **debug ilpower** privileged EXEC command

Disabled Port Caused by False Link Up

If a Cisco powered device is connected to a port and you configure the port by using the **power inline never** interface configuration command, a false link up can occur, placing the port into an error-disabled state. To take the port out of the error-disabled state, enter the **shutdown** and the **no shutdown** interface configuration commands.

You should not connect a Cisco powered device to a port that has been configured with the **power inline never** command.

SFP Module Security and Identification

Cisco 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, see the system message guide for this release.

If you are using a non-Cisco SFP module, remove the SFP module from the switch, and replace it with a Cisco module. After inserting a Cisco 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, see 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.

Monitoring SFP Module Status

You can check the physical or operational status of an SFP module by using the **show interfaces transceiver** privileged EXEC command. This command shows the operational status, such as the temperature and the current for an SFP module on a specific interface and the alarm status. You can also use the command to check the speed and the duplex settings on an SFP module. For more information, see the **show interfaces transceiver** command in the command reference for this release.

Using Ping

- [Understanding Ping, page 1-14](#)
- [Executing Ping, page 1-14](#)

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

Beginning in privileged EXEC mode, use this command to ping another device on the network from the switch:

Command	Purpose
ping ip <i>host address</i>	Ping a remote host through IP or by supplying the hostname 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 1-1 describes the possible ping character output.

Table 1-1 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 end a ping session, enter the escape sequence (**Ctrl-^ X** by default). Simultaneously press and release the **Ctrl**, **Shift**, and **6** keys and then press the **X** key.

Using Layer 2 Traceroute

- [Understanding Layer 2 Traceroute, page 1-15](#)
- [Usage Guidelines, page 1-16](#)
- [Displaying the Physical Path, page 1-17](#)

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 finds 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.
For a list of switches that support Layer 2 traceroute, see the [“Usage Guidelines” section on page 1-16](#). If any devices in the physical path are transparent to CDP, the switch cannot identify the path through these devices. For more information about enabling CDP, see [Chapter 1, “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:

- **tracetroute mac** [**interface** *interface-id*] {*source-mac-address*} [**interface** *interface-id*] {*destination-mac-address*} [**vlan** *vlan-id*] [**detail**]
- **tracetroute mac ip** {*source-ip-address* | *source-hostname*} {*destination-ip-address* | *destination-hostname*} [**detail**]

For more information, see the command reference for this release.

Using IP Traceroute

- [Understanding IP Traceroute, page 1-17](#)
- [Executing IP Traceroute, page 1-18](#)

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 **tracetroute** privileged EXEC command and might or might not appear as a hop in the **tracetroute** 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 **tracetroute** 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 an Internet Control Message Protocol (ICMP) time-to-live-exceeded message to the sender. Traceroute finds 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 learn 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 that this message was sent by the destination port.

Executing IP Traceroute

Beginning in privileged EXEC mode, follow this step to trace that the path packets take through the network:

Command	Purpose
<code>traceroute ip host</code>	Trace the path that packets take through the network.



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

 0 172.2.52.1 0 msec 0 msec 4 msec
 1 172.2.1.203 12 msec 8 msec 0 msec
 2 171.9.16.6 4 msec 0 msec 0 msec
 3 171.9.4.5 0 msec 4 msec 0 msec
 4 171.9.121.34 0 msec 4 msec 4 msec
 5 171.9.15.9 120 msec 132 msec 128 msec
 6 171.9.15.10 132 msec 128 msec 128 msec
Switch#
```

The display shows the hop count, the IP address of the router, and the round-trip time in milliseconds for each of the three probes that are sent.

Table 1-2 Traceroute Output Display Characters

Character	Description
*	The probe timed out.
?	Unknown packet type.
A	Administratively unreachable. Usually, this output means that an access list is blocking traffic.
H	Host unreachable.
N	Network unreachable.
P	Protocol unreachable.
Q	Source quench.
U	Port unreachable.

To end a trace in progress, enter the escape sequence (**Ctrl-^ X** by default). Simultaneously press and release the **Ctrl**, **Shift**, and **6** keys and then press the **X** key.

Using TDR

- [Understanding TDR, page 1-19](#)
- [Running TDR and Displaying the Results, page 1-19](#)

Understanding TDR

You can use the Time Domain Reflector (TDR) feature to diagnose and resolve cabling problems. When running TDR, a local device sends a signal through a cable and compares the reflected signal to the initial signal.

TDR can detect these cabling problems:

- Open, broken, or cut twisted-pair wires—The wires are not connected to the wires from the remote device.
- Shorted twisted-pair wires—The wires are touching each other or the wires from the remote device. For example, a shorted twisted pair can occur if one wire of the twisted pair is soldered to the other wire.

If one of the twisted-pair wires is open, TDR can find the length at which the wire is open.

Use TDR to diagnose and resolve cabling problems in these situations:

- Replacing a switch
- Setting up a wiring closet
- Troubleshooting a connection between two devices when a link cannot be established or when it is not operating properly

Running TDR and Displaying the Results

When you run TDR on an interface, you can run it on the stack master or a stack member.



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

To run TDR, enter the **test cable-diagnostics tdr interface** *interface-id* privileged EXEC command:

To display the results, enter the **show cable-diagnostics tdr interface** *interface-id* privileged EXEC command. For a description of the fields in the display, see the command reference for this release.

Using Debug Commands

These sections explain how you use **debug** commands to diagnose and resolve internetworking problems:

- [Enabling Debugging on a Specific Feature, page 1-20](#)
- [Enabling All-System Diagnostics, page 1-20](#)
- [Redirecting Debug and Error Message Output, page 1-21](#)

**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, see the command reference for this release.

Enabling Debugging on a Specific Feature

When you enable debugging, it is enabled only on the stack master. To enable debugging on a stack member, you must start a session from the stack master by using the **session switch-number** privileged EXEC command. Then, enter the **debug** command at the command-line prompt of the stack member.

**Note**

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

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.

When stack members generate a system error message, the stack master displays the error message to all stack members. The syslog resides on the stack master.

**Note**

Make sure to save the syslog to flash memory so that the syslog is not lost if the stack master fails.

For more information about system message logging, see [Chapter 1, “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, see 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 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 gigabitethernet0/1 vlan 5 1.1.1 2.2.2 ip 13.1.1.1 13.2.2.2
udp 10 20
Global Port Number:24, Asic Number:5
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_00050002_00020002-00_00000000_00000000    00C71  0000002B
Station Descriptor:02340000, DestIndex:0239, RewriteIndex:F005

=====
Egress:Asic 2, 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
Gi0/1     0005     0001.0001.0001  0002.0002.0002

-----
Packet 2
  Lookup                Key-Used                Index-Hit  A-Data
OutptACL 50_0D020202_0D010101-00_40000014_000A0000    01FFE  03000000

Port      Vlan      SrcMac          DstMac          Cos  Dscpv
Gi0/2     0005     0001.0001.0001  0002.0002.0002

-----
<output truncated>
-----
Packet 10
  Lookup                Key-Used                Index-Hit  A-Data
OutptACL 50_0D020202_0D010101-00_40000014_000A0000    01FFE  03000000
Packet dropped due to failed DEJA_VU Check on Gi1/0/2
Packet dropped due to failed DEJA_VU Check on Gi0/2
```

This is an example of the output when the packet coming in on 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 gigabitethernet0/1 vlan 5 1.1.1 0009.43a8.0145 ip 13.1.1.1
13.2.2.2 udp 10 20
Global Port Number:24, Asic Number:5
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
OutputACL 50_0D020202_0D010101-00_40000014_000A0000    01FFE    03000000

Port                Vlan      SrcMac                DstMac      Cos  Dscp
interface-id        0005     0001.0001.0001      0009.43A8.0145

```

Using the crashinfo Files

The crashinfo files save 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. The switch creates two types of crashinfo files:

- Basic crashinfo file—The switch automatically creates this file the next time you boot up the Cisco IOS image after the failure.
- Extended crashinfo file—The switch automatically creates this file when the system is failing.

Basic crashinfo Files

The information in the basic file includes the Cisco IOS image name and version that failed, a list of the processor registers, and other switch-specific information. You can provide this information to the Cisco technical support representative by using the **show tech-support** privileged EXEC command.

Basic crashinfo files are kept in this directory on the flash file system:

```
flash:/crashinfo/.
```

The filenames are `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 basic 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.

Extended crashinfo Files

The switch creates the extended crashinfo file when the system is failing. The information in the extended file includes additional information that can help determine the cause of the switch failure. You provide this information to the Cisco technical support representative by manually accessing the file and using the **more** or the **copy** privileged EXEC command.

Extended crashinfo files are kept in this directory on the flash file system:

```
flash:/crashinfo_ext/.
```

The filenames are `crashinfo_ext_n` where *n* is a sequence number.

You can configure the switch to not create the extended crashinfo file by using the **no exception crashinfo** global configuration command.

Using On-Board Failure Logging



Note

OBFL is supported only on Catalyst 2960-S switches running the LAN base image.

You can use the on-board-failure logging (OBFL) feature to collect information about the switch. The information includes uptime, temperature, and voltage information and helps Cisco technical support representatives to troubleshoot switch problems. We recommend that you keep OBFL enabled and do not erase the data stored in the flash memory.

This section has this information:

- [Understanding OBFL, page 1-24](#)
- [Configuring OBFL, page 1-25](#)
- [Displaying OBFL Information, page 1-25](#)

Understanding OBFL

By default, OBFL is enabled. It collects information about the switch and small form-factor pluggable (SFP) modules. The switch stores this information in the flash memory:

- CLI commands—Record of the OBFL CLI commands that are entered on a standalone switch or a switch stack member
- Environment data—Unique device identifier (UDI) information for a standalone switch or a stack member and for all the connected FRU devices: the product identification (PID), the version identification (VID), and the serial number
- Message—Record of the hardware-related system messages generated by a standalone switch or a stack member
- Power over Ethernet (PoE)—Record of the power consumption of PoE ports on a standalone switch or a stack member
- Temperature—Temperature of a standalone switch or a stack member
- Uptime data—Time when a standalone switch or a stack member starts, the reason the switch restarts, and the length of time the switch has been running since it last restarted
- Voltage—System voltages of a standalone switch or a stack member

You should manually set the system clock or configure it by using Network Time Protocol (NTP).

When the switch is running, you can retrieve the OBFL data by using the **show logging onboard** privileged EXEC commands. If the switch fails, contact your Cisco technical support representative to find out how to retrieve the data.

When an OBFL-enabled switch is restarted, there is a 10-minute delay before logging of new data begins.



Note

OBFL is supported *only* on Catalyst 2960-S switches. It is not supported on Catalyst 2960 switches.

Configuring OBFL

To enable OBFL, use the **hw-module module** *[switch-number]* **logging onboard** **[message level level]** global configuration command. On Catalyst 2960-S switches, the range for *switch-number* is from 1 to 4. Use the **message level level** parameter to specify the severity of the hardware-related messages that the switch generates and stores in the flash memory.

To copy the OBFL data to the local network or a specific file system, use the **copy logging onboard module stack-member destination** privileged EXEC command.



Caution

We recommend that you do not disable OBFL and that you do not remove the data stored in the flash memory.

To disable OBFL, use the **no hw-module module** *[switch-number]* **logging onboard** **[message level]** global configuration command.

To clear all the OBFL data in the flash memory except for the uptime and CLI command information, use the **clear logging onboard** privileged EXEC command.

In a switch stack, you can enable OBFL on a standalone switch or on all stack members by using the **hw-module module logging onboard** **[message level level]** global configuration command.

For more information about the commands in this section, see the command reference for this release.

Displaying OBFL Information

To display the OBFL information, use one or more of the privileged EXEC commands in [Table 1-3](#):

Table 1-3 *Commands for Displaying OBFL Information*

Command	Purpose
show logging onboard [module <i>[switch-number]</i>]] cilog	Displays the OBFL CLI commands that were entered on a standalone switch or the specified stack members.
show logging onboard [module <i>[switch-number]</i>]] environment	Display the UDI information for a standalone switch or the specified stack members and for all the connected FRU devices: the PID, the VID, and the serial number.
show logging onboard [module <i>[switch-number]</i>]] message	Display the hardware-related messages generated by a standalone switch or the specified stack members.
show logging onboard [module <i>[switch-number]</i>]] poe	Display the power consumption of PoE ports on a standalone switch or the specified stack members.
show logging onboard [module <i>[switch-number]</i>]] temperature	Display the temperature of a standalone switch or the specified switch stack members.
show logging onboard [module <i>[switch-number]</i>]] uptime	Display the time when a standalone switch or the specified stack members start, the reason the standalone switch or specified stack members restart, and the length of time that the standalone switch or specified stack members have been running since they last restarted.
show logging onboard [module <i>[switch-number]</i>]] voltage	Display the system voltages of a standalone switch or the specified stack members.

For more information about using the commands in [Table 1-3](#) and for examples of OBFL data, see the command reference for this release.

Memory Consistency Check Routines

The switch runs memory consistency check routines to detect and correct invalid ternary content addressable memory (TCAM) table entries that can affect the performance of the switch.

If the switch cannot fix the error, it logs a system error message, specifying the TCAM space in which the error is located:

- Unassigned space: Unassigned TCAM table entries for the current SDM template.



Note Unassigned spaces do not apply to 2960-S switches.

- Hult Forwarding TCAM Manager (HFTM) space: Related to the Layer 2 and Layer 3 forwarding tables.
- Hult Quality of Service (QoS)/access control list (ACL) TCAM Manager (HQATM) space: Related to ACL and ACL-like tables such as QoS classification and policy routing.

The output from the **show platform tcam errors** privileged EXEC command provides information about the TCAM memory consistency integrity on the switch.

Beginning in privileged EXEC mode, use the **show platform tcam errors** command to display the TCAM memory consistency check errors detected on the switch:

Command	Purpose
show platform tcam errors	Displays TCAM memory consistency check errors in the HQATM, HFTM, and unassigned spaces on the TCAM.

This example shows the output of the **show platform tcam errors** command:

```
DomainMember# show platform tcam errors

TCAM Memory Consistency Checker Errors
-----
TCAM Space Values  Masks  Fixups  Retries Failures
Unassigned  0      0      0      0      0
HFTM        0      0      0      0      0
HQATM       0      0      0      0      0
```

DomainMember#



Note For 2960-S switch output unassigned spaces do not apply.

Table 1-4 Definitions of Fields in TCAM Checker Output

Column	Description
Values	The number of invalid values found in the TCAM tables.
Masks	The number of invalid masks found in the TCAM tables.

Table 1-4 Definitions of Fields in TCAM Checker Output (continued)

Column	Description
Fixups	The number of initial attempts to fix the invalid values or masks.
Retries	The number of attempts to fix the invalid values or masks.
Failures	The number of failed attempts to fix the invalid values or masks.

For more information about the **show platform tcam errors** privileged EXEC command, see the command reference for this release.

Troubleshooting Tables

These tables are a condensed version of troubleshooting documents on Cisco.com.

- [“Troubleshooting CPU Utilization” section on page 1-27](#)
- [“Troubleshooting Power over Ethernet \(PoE\)” section on page 1-28](#)
- [“Troubleshooting Switch Stacks” section on page 1-32](#)

Troubleshooting CPU Utilization

This section lists some possible symptoms that could be caused by the CPU being too busy and shows how to verify a CPU utilization problem. [Table 1-5](#) lists the primary types of CPU utilization problems that you can identify. It gives possible causes and corrective action with links to the [Troubleshooting High CPU Utilization](#) document on Cisco.com.

Possible Symptoms of High CPU Utilization

Note that excessive CPU utilization might result in these symptoms, but the symptoms could also result from other causes.

- Spanning tree topology changes
- EtherChannel links brought down due to loss of communication
- Failure to respond to management requests (ICMP ping, SNMP timeouts, slow Telnet or SSH sessions)
- UDLD flapping
- IP SLAs failures because of SLAs responses beyond an acceptable threshold
- DHCP or IEEE 802.1x failures if the switch does not forward or respond to requests

Layer 3 switches:

- Dropped packets or increased latency for packets routed in software
- BGP or OSPF routing topology changes
- HSRP flapping

Verifying the Problem and Cause

To determine if high CPU utilization is a problem, enter the **show processes cpu sorted** privileged EXEC command. Note the underlined information in the first line of the output example.

```
Switch# show processes cpu sorted
CPU utilization for five seconds: 8%/0%; one minute: 7%; five minutes: 8%
PID Runtime(ms) Invoked uSecs 5Sec 1Min 5Min TTY Process
309 42289103 752750 56180 1.75% 1.20% 1.22% 0 RIP Timers
140 8820183 4942081 1784 0.63% 0.37% 0.30% 0 HRPC qos request
100 3427318 16150534 212 0.47% 0.14% 0.11% 0 HRPC pm-counters
192 3093252 14081112 219 0.31% 0.14% 0.11% 0 Spanning Tree
143 8 37 216 0.15% 0.01% 0.00% 0 Exec
...
<output truncated>
```

This example shows normal CPU utilization. The output shows that utilization for the last 5 seconds is 8%/0%, which has this meaning:

- The total CPU utilization is 8 percent, including both time running Cisco IOS processes and time spent handling interrupts
- The time spent handling interrupts is zero percent.

Table 1-5 Troubleshooting CPU Utilization Problems

Type of Problem	Cause	Corrective Action
Interrupt percentage value is almost as high as total CPU utilization value.	The CPU is receiving too many packets from the network.	Determine the source of the network packet. Stop the flow, or change the switch configuration. See the section on “Analyzing Network Traffic.”
Total CPU utilization is greater than 50% with minimal time spent on interrupts.	One or more Cisco IOS process is consuming too much CPU time. This is usually triggered by an event that activated the process.	Identify the unusual event, and troubleshoot the root cause. See the section on “Debugging Active Processes.”

For complete information about CPU utilization and how to troubleshoot utilization problems, see the [Troubleshooting High CPU Utilization](#) document on Cisco.com.

Troubleshooting Power over Ethernet (PoE)

[Table 1-6](#) lists some PoE troubleshooting scenarios. For more information causes and solutions referenced in the table, see the [Troubleshooting Power over Ethernet \(PoE\)](#) troubleshooting guide on Cisco.com.



Note

Power over Ethernet Plus (PoE+) is not supported on Catalyst 2960-S switches.

Table 1-6 Power Over Ethernet Troubleshooting Scenarios

Symptom or problem	Possible cause and solution
<p>No PoE on only one port.</p> <p>Trouble is on only one switch port. PoE and non-PoE devices do not work on this port, but do on other ports.</p>	<p>Verify that the powered device works on another PoE port.</p> <p>Use the show run, show interface status, or show power inline detail user EXEC commands to verify that the port is not shut down or error disabled.</p> <p>Note Most switches turn off port power when the port is shut down, even though the IEEE specifications make this optional.</p> <p>Verify that the Ethernet cable from the powered device to the switch port is good: Connect a known good non-PoE Ethernet device to the Ethernet cable, and make sure that the powered device establishes a link and exchanges traffic with another host.</p> <p>Verify that the total cable length from the switch front panel to the powered device is not more than 100 meters.</p> <p>Disconnect the Ethernet cable from the switch port. Use a short Ethernet cable to connect a known good Ethernet device directly to this port on the switch front panel (not on a patch panel). Verify that it can establish an Ethernet link and exchange traffic with another host, or ping the port VLAN SVI. Next, connect a powered device to this port, and verify that it powers on.</p> <p>If a powered device does not power on when connected with a patch cord to the switch port, compare the total number of connected powered devices to the switch power budget (available PoE). Use the show inline power and show inline power detail commands to verify the amount of available power.</p>

Table 1-6 Power Over Ethernet Troubleshooting Scenarios (continued)

Symptom or problem	Possible cause and solution
<p>No PoE on all ports or a group of ports. Trouble is on all switch ports. Nonpowered Ethernet devices cannot establish an Ethernet link on any port, and PoE devices do not power on.</p>	<p>If there is a continuous, intermittent, or reoccurring alarm related to power, replace the power supply if possible it is a field-replacable unit. Otherwise, replace the switch.</p> <p>If the problem is on a consecutive group of ports but not all ports, the power supply is probably not defective, and the problem could be related to PoE regulators in the switch.</p> <p>Use the show log privileged EXEC command to review alarms or system messages that previously reported PoE conditions or status changes.</p> <p>If there are no alarms, use the show interface status command to verify that the ports are not shut down or error-disabled. If ports are error-disabled, use the shut and no shut interface configuration commands to re-enable the ports.</p> <p>Use the show env power and show power inline privileged EXEC commands to review the PoE status and power budget (available PoE).</p> <p>Review the running configuration to verify that power inline never is not configured on the ports.</p> <p>Connect a nonpowered Ethernet device directly to a switch port. Use only a short patch cord. Do not use the existing distribution cables. Enter the shut and no shut interface configuration commands, and verify that an Ethernet link is established. If this connection is good, use a short patch cord to connect a powered device to this port and verify that it powers on. If the device powers on, verify that all intermediate patch panels are correctly connected.</p> <p>Disconnect all but one of the Ethernet cables from switch ports. Using a short patch cord, connect a powered device to only one PoE port. Verify the powered device does not require more power than can be delivered by the switch port.</p> <p>Use the show power inline privileged EXEC command to verify that the powered device can receive power when the port is not shut down. Alternatively, watch the powered device to verify that it powers on.</p> <p>If a powered device can power on when only one powered device is connected to the switch, enter the shut and no shut interface configuration commands on the remaining ports, and then reconnect the Ethernet cables one at a time to the switch PoE ports. Use the show interface status and show power inline privileged EXEC commands to monitor inline power statistics and port status.</p> <p>If there is still no PoE at any port, a fuse might be open in the PoE section of the power supply. This normally produces an alarm. Check the log again for alarms reported earlier by system messages.</p>

Table 1-6 Power Over Ethernet Troubleshooting Scenarios (continued)

Symptom or problem	Possible cause and solution
<p>Cisco IP Phone disconnects or resets.</p> <p>After working normally, a Cisco phone or wireless access point intermittently reloads or disconnects from PoE.</p>	<p>Verify all electrical connections from the switch to the powered device. Any unreliable connection results in power interruptions and irregular powered device functioning such as erratic powered device disconnects and reloads.</p> <p>Verify that the cable length is not more than 100 meters from the switch port to the powered device.</p> <p>Notice what changes in the electrical environment at the switch location or what happens at the powered device when the disconnect occurs?</p> <p>Notice whether any error messages appear at the same time a disconnect occurs. Use the show log privileged EXEC command to review error messages.</p> <p>Verify that an IP phone is not losing access to the Call Manager immediately before the reload occurs. (It might be a network problem and not a PoE problem.)</p> <p>Replace the powered device with a non-PoE device, and verify that the device works correctly. If a non-PoE device has link problems or a high error rate, the problem might be an unreliable cable connection between the switch port and the powered device.</p>
<p>Non-Cisco powered device does not work on Cisco PoE switch.</p> <p>A non-Cisco powered device is connected to a Cisco PoE switch, but never powers on or powers on and then quickly powers off. Non-PoE devices work normally.</p>	<p>Use the show power inline command to verify that the switch power budget (available PoE) is not depleted before or after the powered device is connected. Verify that sufficient power is available for the powered device type before you connect it.</p> <p>Use the show interface status command to verify that the switch detects the connected powered device.</p> <p>Use the show log command to review system messages that reported an overcurrent condition on the port. Identify the symptom precisely: Does the powered device initially power on, but then disconnect? If so, the problem might be an initial surge-in (or <i>inrush</i>) current that exceeds a current-limit threshold for the port.</p>

Troubleshooting Switch Stacks

Table 1-7 lists some switch stack troubleshooting scenarios. For more detailed information about causes and solutions referenced in the table, set the *Troubleshooting Switch Stacks* guide on Cisco.com.



Note Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

Table 1-7 Switch Stack Troubleshooting Scenarios

Symptom/problem	How to Verify Problem	Possible Cause/Solution
General troubleshooting of switch stack issues	Review this document.	Use the <i>Troubleshooting Switch Stacks</i> document for problem solutions and tutorial information.
Switch cannot join stack	Enter the show switch privileged EXEC command.	Incompatible Cisco IOS versions between stack members and new switch.
	Enter the show version user EXEC command.	Incompatible license levels in a Catalyst 3750-E switch..
	Enter the show platform stack-manager all command.	Incompatible Cisco IOS version numbers between stack members and new switch..
	Look carefully at the cables and connections.	Unreliable StackWise cable or incomplete connection.
	Enter the show sdm prefer command.	Configuration mismatch (that is, SDM templates) if switch was used for other applications before you added it to the stack. Incompatible IOS version between stack members and new switch .
StackWise port frequently or rapidly changing up/down states (flapping)	Error messages report stack link problems. Possible traffic disruption.	Unreliable StackWise cable connection or interface.
Switch member port not coming up	Enter the show switch detail privileged EXEC command.	Unreliable StackWise cable connection or interface.
Reduced stack ring bandwidth, or slow throughput between switch ports or between switches in the stack.	Enter the show switch stack-ring speed user EXEC command.	Bad connection between StackWise cable connection and switch chassis connector .
	Enter the show switch detail user EXEC command to see which stack cable or connection is causing the problem.	Defective or missing StackWise cable .
	<ul style="list-style-type: none"> • Check the retainer screws on the StackWise cable connectors. • Enter the show switch privileged EXEC command to see whether new switch shows as Ready, Progressing, or Provisioned. 	<ul style="list-style-type: none"> • Loose retainer screws or overly tightened retainer screws. • Check status of stack members.
Port numbering in one or more switches is incorrect or changed.	Enter the show switch detail user EXEC command.	Multiple StackWise cables are disconnected from stack members creating two separate stacks. .

Table 1-7 Switch Stack Troubleshooting Scenarios (continued)

Symptom/problem	How to Verify Problem	Possible Cause/Solution
Slow traffic throughput on stack ring	Test the switch interface.	Defective StackWise switch interface. Note The only solution is to replace the switch.
Problems with stack master election, stacks merging, or new switches joining stack	Review the rules of stack master election.	Current stack master is rebooted or disconnected.
	Port numbering seems off.	Verify port numbering.
	Enter the show switch privileged EXEC command.	Interpret state messages..
Stack members need to be upgraded.	Stack members running different major or minor versions of the Cisco IOS software.	Defective StackWise switch interface or cable .
StackWise link connection problems	Look at the LED behavior.	Stack not operating at full bandwidth.



CHAPTER 1

Configuring Online Diagnostics

This chapter describes how to configure the online diagnostics on the 2960 and 2960-S switches.



Note

Online Diagnostics is supported only on Catalyst 2960-S switches running the LAN base image.



Note

For complete syntax and usage information for the commands used in this chapter, see the switch command reference for this release.

- [Understanding How Online Diagnostics Work, page 1-1](#)
- [Running Online Diagnostic Tests, page 1-3](#)

Understanding How Online Diagnostics Work

With online diagnostics, you can test and verify the hardware functionality of the switch while the switch is connected to a live network.

The online diagnostics contain packet switching tests that check different hardware components and verify the data path and the control signals.

The online diagnostics detect problems in these areas:

- Hardware components
- Interfaces (Ethernet ports and so forth)
- Solder joints

Online diagnostics are categorized as on-demand, scheduled, or health-monitoring diagnostics. On-demand diagnostics run from the CLI; scheduled diagnostics run at user-designated intervals or at specified times when the switch is connected to a live network; and health-monitoring runs in the background.

Scheduling Online Diagnostics

You can schedule online diagnostics to run at a designated time of day or on a daily, weekly, or monthly basis for a specific switch. Use the **no** form of this command to remove the scheduling.

Beginning in global configuration mode, use this command to schedule online diagnostics:

Command	Purpose
diagnostic schedule switch <i>numtest</i> { <i>test_id</i> <i>test_id_range</i> all basic non-disruptive } { daily <i>hh:mm</i> on <i>mm dd yyyy</i> <i>hh:mm</i> } weekly <i>day_of_week</i> <i>hh:mm</i> }	Schedule on-demand diagnostic tests for a specific date and time, how many times to run the test (iterations), and what action to take when errors are found.

This example shows how to schedule diagnostic testing on a specific date and time for a specific switch:

```
Switch(config)# diagnostic schedule switch 1 test 1,2,4-6 on january 3 2006 23:32
```

This example shows how to schedule diagnostic testing to occur weekly at a certain time for a specific switch:

```
Switch(config)# diagnostic schedule switch 1 test 1,2,4-6 weekly friday 09:23
```

Configuring Health-Monitoring Diagnostics

You can configure health-monitoring diagnostic testing on a specified switch while the switch is connected to a live network. You can configure the execution interval for each health-monitoring test, whether or not to generate a system message upon a test failure, or to enable or disable an individual test. Use the **no** form of this command to disable testing.

Beginning in global configuration mode, use these commands to configure health-monitoring diagnostics:

Command	Purpose
diagnostic monitor interval switch <i>numtest</i> { <i>test_id</i> <i>test_id_range</i> } <i>hour:mm:ss</i> <i>milliseconds</i> <i>day</i>	Configure the health-monitoring interval of the specified tests for the specified switch. By default, monitoring is disabled.
diagnostic monitor syslog	Enable the generation of a syslog message for health-monitoring test failures. By default, syslog is disabled.
diagnostic monitor threshold switch <i>numtest</i> { <i>test_id</i> <i>test_id_range</i> } failure count <i>count</i>	Set the failure threshold for monitoring tests. By default, monitoring is disabled.

Use the **no diagnostic monitor interval switch** {*num*} **test** {*test-id* | *test-id-range* | **all**} global configuration command to change the interval to the default value or to zero. Use the **no diagnostic monitor syslog** command to disable generation of syslog messages when a health-monitoring test fails. Use the **diagnostic monitor threshold switch** *numtest* {*test_id* | *test_id_range* | **all**} **failure count** command to remove the failure threshold.

This example shows how to configure the specified test to run every 2 minutes:

```
Switch(config)# diagnostic monitor interval switch 1 test 1 00:02:00 0 1
```

This example shows how to set the failure threshold for test monitoring on a switch:

```
Switch(config)# diagnostic monitor threshold switch 1 test 1 failure count 50
```

This example shows how to enable the generation of a syslog message when any health monitoring test fails:

```
Switch(config)# diagnostic monitor syslog
```

Running Online Diagnostic Tests

After you configure online diagnostics, you can start diagnostic tests or display the test results. You can also see which tests are configured for each switch and what diagnostic tests have already run.

These sections describe how to run online diagnostic tests after they have been configured:

- [Starting Online Diagnostic Tests, page 1-3](#)
- [Displaying Online Diagnostic Tests and Test Results, page 1-4](#)

Starting Online Diagnostic Tests

After you configure diagnostic tests to run on the switch or on individual switches, you can use **start** to begin a diagnostic test.

Beginning in global configuration mode, use this command to start an online diagnostic test:

Command	Purpose
diagnostic start switch <i>num</i> test { <i>test-id</i> <i>test-id-range</i> all basic non-disruptive }	Start a diagnostic test on a specific switch.

This example shows how to start a diagnostic test on a specific switch:

```
Switch# diagnostic start switch 1 test 1
Switch#
06:27:50: %DIAG-6-TEST_RUNNING: Switch 1: Running TestPortAsicStackPortLoopback{ID=1} ...
(switch-1)
06:27:51: %DIAG-6-TEST_OK: Switch 1: TestPortAsicStackPortLoopback{ID=1} has completed
successfully (switch-1)
Switch#
```

This example shows how to start diagnostics test 2 on a switch disrupting normal system operations, causing the switch to lose stack connectivity, and then to reload:

```
Switch# diagnostic start switch 1 test 2
Switch 1: Running test(s) 2 will cause the switch under test to reload after completion of
the test list.
Switch 1: Running test(s) 2 may disrupt normal system operation
Do you want to continue? [no]: y
Switch#
16:43:29: %STACKMGR-4-STACK_LINK_CHANGE: Stack Port 2 Switch 2 has changed to state DOWN
16:43:30: %STACKMGR-4-STACK_LINK_CHANGE: Stack Port 1 Switch 9 has changed to state DOWN
16:43:30: %STACKMGR-4-SWITCH_REMOVED: Switch 1 has been REMOVED from the stack
Switch#
16:44:35: %STACKMGR-4-STACK_LINK_CHANGE: Stack Port 1 Switch 2 has changed to state UP
16:44:37: %STACKMGR-4-STACK_LINK_CHANGE: Stack Port 2 Switch 2 has changed to state UP
16:44:45: %STACKMGR-4-SWITCH_ADDED: Switch 1 has been ADDED to the stack
16:45:00: %STACKMGR-5-SWITCH_READY: Switch 1 is READY
16:45:00: %STACKMGR-4-STACK_LINK_CHANGE: Stack Port 1 Switch 1 has changed to state UP
16:45:00: %STACKMGR-4-STACK_LINK_CHANGE: Stack Port 2 Switch 1 has changed to state UP
```

```

00:00:20: %STACKMGR-4-SWITCH_ADDED: Switch 1 has been ADDED to the stack (Switch-1)
00:00:20: %STACKMGR-4-SWITCH_ADDED: Switch 2 has been ADDED to the stack (Switch-1)
00:00:25: %SPANTREE-5-EXTENDED_SYSID: Extended SysId enabled for type vlan (Switch-1)
00:00:29: %SYS-5-CONFIG_I: Configured from memory by console (Switch-1)
00:00:29: %STACKMGR-5-SWITCH_READY: Switch 2 is READY (Switch-1)
00:00:29: %STACKMGR-5-MASTER_READY: Master Switch 2 is READY (Switch-1)
00:00:30: %STACKMGR-5-SWITCH_READY: Switch 1 is READY (Switch-1)
00:00:30: %DIAG-6-TEST_RUNNING: Switch 1: Running TestPortAsicLoopback{ID=2} ...
(Switch-1)
00:00:30: %DIAG-6-TEST_OK: Switch 1: TestPortAsicLoopback{ID=2} has completed successfully
(Switch-1)

```

You see this message if the test will cause a stack partition:

```

Switch 6: Running test(s) 2 will cause the switch under test to reload after completion of
the test list.
Switch 6: Running test(s) 2 will partition stack
Switch 6: Running test(s) 2 may disrupt normal system operation
Do you want to continue? [no]:

```

Displaying Online Diagnostic Tests and Test Results

You can display the online diagnostic tests that are configured for specific switches and check the results of the tests using the **show** commands.

To display the diagnostic tests that are configured for a switch and the test results, use these privileged EXEC commands:

Table 1-1 show diagnostic Commands

Command	Purpose
show diagnostic content switch [<i>num</i> all]	Display the online diagnostics configured for a switch.
show diagnostic status	Display whether a switch is running a test.
show diagnostic result switch [<i>num</i> all] detail	Display the online diagnostics test results.
show diagnostic result switch [<i>num</i> all] test [<i>test_id</i> <i>test_id_range</i> all] [detail]	
show diagnostic schedule switch [<i>num</i> all]	Display the online diagnostics test schedule.
show diagnostic post	Display the results of POST. (The same as the show post command.)

This example shows how to display the online diagnostics that are configured on a switch:

```

Switch# show diagnostic content switch 3
Switch 3:
Diagnostics test suite attributes:
  B/* - Basic ondemand test / NA
  P/V/* - Per port test / Per device test / NA
  D/N/* - Disruptive test / Non-disruptive test / NA
  S/* - Only applicable to standby unit / NA
  X/* - Not a health monitoring test / NA
  F/* - Fixed monitoring interval test / NA
  E/* - Always enabled monitoring test / NA
  A/I - Monitoring is active / Monitoring is inactive
  R/* - Switch will reload after test list completion / NA
  P/* - will partition stack / NA

```

ID	Test Name	attributes	Test Interval day hh:mm:ss.ms	Thre- shold
1)	TestPortAsicStackPortLoopback	B*N***A**	000 00:01:00.00	n/a
2)	TestPortAsicLoopback	B*D*X**IR*	not configured	n/a
3)	TestPortAsicCam	B*D*X**IR*	not configured	n/a
4)	TestPortAsicRingLoopback	B*D*X**IR*	not configured	n/a
5)	TestMicRingLoopback	B*D*X**IR*	not configured	n/a
6)	TestPortAsicMem	B*D*X**IR*	not configured	n/a

This example shows how to display the online diagnostic results for a switch:

```
Switch# show diagnostic result
Switch 1: SerialNo :
Overall diagnostic result: PASS
Test results: (. = Pass, F = Fail, U = Untested)
1) TestPortAsicStackPortLoopback ---> .
2) TestPortAsicLoopback -----> .
3) TestPortAsicCam -----> .
4) TestPortAsicRingLoopback -----> .
5) TestMicRingLoopback -----> .
6) TestPortAsicMem -----> .
```

This example shows how to display the online diagnostic test status:

```
Switch# show diagnostic status
<BU> - Bootup Diagnostics, <HM> - Health Monitoring Diagnostics,
<OD> - OnDemand Diagnostics, <SCH> - Scheduled Diagnostics
=====
Card   Description                               Current Running Test           Run by
-----
1      N/A                                         N/A                             N/A
2      TestPortAsicStackPortLoopback             TestPortAsicLoopback          <OD>
      TestPortAsicCam                           TestPortAsicCam               <OD>
      TestPortAsicRingLoopback                 TestMicRingLoopback           <OD>
      TestPortAsicMem                           N/A                            <OD>
3      N/A                                         N/A                             N/A
4      N/A                                         N/A                             N/A
=====
Switch#
```

This example shows how to display the online diagnostic test schedule for a switch:

```
Switch# show diagnostic schedule switch 1
Current Time = 14:39:49 PST Tue Jul 5 2005
Diagnostic for Switch 1:
Schedule #1:
To be run daily 12:00
Test ID(s) to be executed: 1.
```




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

This appendix describes how to manipulate the Catalyst 2960 and 2960 switch flash file system, how to copy configuration files, and how to archive (upload and download) software images to a switch. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.

**Note**

For complete syntax and usage information for the commands used in this chapter, see the switch command reference for this release and the *Cisco IOS Configuration Fundamentals Command Reference, Release 12.4* on Cisco.com.

This appendix has these sections:

- [Working with the Flash File System, page 1-1](#)
- [Working with Configuration Files, page 1-8](#)
- [Working with Software Images, page 1-24](#)

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:*.

As viewed from the stack master, or any stack member, *flash:* refers to the local flash device, which is the device attached to the same switch on which the file system is being viewed. In a switch stack, each of the flash devices from the various stack members can be viewed from the stack master. The names of these flash file systems include the corresponding switch member numbers. For example, *flash3:*, as viewed from the stack master, refers to the same file system as does *flash:* on stack member 3. Use the **show file systems** privileged EXEC command to list all file systems, including the flash file systems in the switch stack.

No more than one user at a time can manage the software images and configuration files for a switch stack.

These sections contain this configuration information:

- [Displaying Available File Systems, page 1-2](#)
- [Setting the Default File System, page 1-3](#)
- [Displaying Information about Files on a File System, page 1-3](#)

- [Creating and Removing Directories, page 1-4](#)
- [Copying Files, page 1-5](#)
- [Deleting Files, page 1-5](#)
- [Creating, Displaying, and Extracting tar Files, page 1-6](#)
- [Displaying the Contents of a File, page 1-8](#)

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. In this example, the stack master is stack member 3; therefore `flash3:` is aliased to `flash:.` The file system on stack member 5 is displayed as `flash5` on the stack master.

```
Switch# show file systems
File Systems:
  Size(b)    Free(b)    Type  Flags  Prefixes
*  15998976   5135872   flash rw    flash:flash3:
    -        -         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:
    15998976   645120   unknown rw    flash5:
    -        -         network rw    rcp:
    -        -         network rw    ftp:
```

Table 1-1 *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. flash —The file system is for a flash memory device. nvram —The file system is for a NVRAM device. opaque —The file system is a locally generated <i>pseudo</i> file system (for example, the <i>system</i>) or a download interface, such as brimux. unknown —The file system is an unknown type.

Table 1-1 *show file systems Field Descriptions (continued)*

Field	Value
Flags	Permission for file system. ro —read-only. rw —read/write. wo —write-only.
Prefixes	Alias for file system. flash: —Flash file system. nvr: —NVRAM. null: —Null destination for copies. You can copy a remote file to null to find its size. rep: —Remote Copy Protocol (RCP) network server. system: —Contains the system memory, including the running configuration. tftp: —TFTP network server. xmodem: —Obtain the file from a network machine by using the Xmodem protocol. ymodem: —Obtain the file from a network machine by using the Ymodem protocol.

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 1-2](#):

Table 1-2 *Commands for Displaying Information About Files*

Command	Description
dir [/all] [filesystem:][filename]	Display a list of files on a file system.
show file systems	Display more information about each of the files on a file system.

Table 1-2 Commands for Displaying Information About Files (continued)

Command	Description
show file information <i>file-url</i>	Display information about a specific file.
show file descriptors	Display a list of open file descriptors. File descriptors are the internal representations of open files. You can use this command to see if another user has a file open.

Changing Directories and Displaying the Working Directory

Beginning in privileged EXEC mode, follow these steps to change directories and display the working directory.

	Command	Purpose
Step 1	dir <i>filesystem:</i>	Display the directories on the specified file system. For <i>filesystem:</i> , use flash: for the system board flash device.
Step 2	cd <i>new_configs</i>	Change to the directory of interest. The command example shows how to change to the directory named <i>new_configs</i> .
Step 3	pwd	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	dir <i>filesystem:</i>	Display the directories on the specified file system. For <i>filesystem:</i> , use flash: for the system board flash device.
Step 2	mkdir <i>old_configs</i>	Create a new directory. The command example shows how to create the directory named <i>old_configs</i> . Directory names are case sensitive. Directory names are limited to 45 characters between the slashes (/); the name cannot contain control characters, spaces, deletes, slashes, quotes, semicolons, or colons.
Step 3	dir <i>filesystem:</i>	Verify your entry.

To delete a directory with all its files and subdirectories, use the **delete /force /recursive** *filesystem:/file-url* privileged EXEC command.

Use the **/recursive** keyword to delete the named directory and all subdirectories and the files contained in it. Use the **/force** keyword to suppress the prompting that confirms a deletion of each file in the directory. You are prompted only once at the beginning of this deletion process. Use the **/force** and **/recursive** keywords for deleting old software images that were installed by using the **archive download-sw** command but are no longer needed.

For *filesystem*, use **flash:** for the system board flash device. For *file-url*, enter the name of the directory to be deleted. All the files in the directory and the directory are removed.

**Caution**

When files and directories are deleted, their contents cannot be recovered.

Copying Files

To copy a file from a source to a destination, use the **copy 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 1-8.

To copy software images either by downloading a new version or by 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 1-24.

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 **rnp:[[/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-name.tar
image-name/ (directory)
image-name/html/ (directory)
image-name/html/foo.html (0 bytes)
image-name/image-name.bin (610856 bytes)
image-name/info (219 bytes)
```

This example shows how to display only the */html* directory and its contents:

```
Switch# archive tar /table flash: image-name/html
cimage-name/html
cimage-name/html/ (directory)
cimage-name/html/const.htm (556 bytes)
cimage-name/html/xhome.htm (9373 bytes)
cimage-name/html/menu.css (1654 bytes)
<output truncated>
```

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 | /ebdic] *file-url* privileged EXEC command:

This example shows how to display the contents of a configuration file on a TFTP server:

```
Switch# more tftp://serverA/hampton/savedconfig
!
! Saved configuration on server
!
version 11.3
service timestamps log datetime localtime
service linenumber
service udp-small-servers
service pt-vty-logging
!
<output truncated>
```

Working with Configuration Files

This section describes how to create, load, and maintain configuration files.



Note

For information about configuration files in switch stacks, see the [“Stack Configuration Files” section on page 1-13](#).

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 1, “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.

These sections contain configuration information:

- [Guidelines for Creating and Using Configuration Files](#), page 1-9
- [Configuration File Types and Location n](#), page 1-10
- [Creating a Configuration File By Using a Text Editor](#), page 1-10
- [Copying Configuration Files By Using TFTP](#), page 1-10
- [Copying Configuration Files By Using FTP](#), page 1-12
- [Copying Configuration Files By Using RCP](#), page 1-16
- [Clearing Configuration Information](#), page 1-19
- [Replacing and Rolling Back Configurations](#), page 1-19

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 1-11, the [“Downloading a Configuration File By Using FTP”](#) section on page 1-13, or the [“Downloading a Configuration File By Using RCP”](#) section on page 1-17.
- 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.

These sections contain this configuration information:

- [Preparing to Download or Upload a Configuration File By Using TFTP, page 1-10](#)
- [Downloading the Configuration File By Using TFTP, page 1-11](#)
- [Uploading the Configuration File By Using TFTP, page 1-12](#)

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, see 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 1-10.
 - 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 hostname of the TFTP server and the name of the file to download.

Use one of these privileged EXEC commands:

- **copy tftp:[[/location]/directory]/filename system:running-config**
- **copy tftp:[[/location]/directory]/filename nvram:startup-config**

The configuration file downloads, and the commands are executed as the file is parsed line-by-line.

This example shows how to configure the software from the file `tokyo-config` at IP address 172.16.2.155:

```
Switch# copy tftp://172.16.2.155/tokyo-config system:running-config
Configure using tokyo-config from 172.16.2.155? [confirm] y
Booting tokyo-config from 172.16.2.155:!!! [OK - 874/16000 bytes]
```

Uploading the Configuration File By Using TFTP

To upload a configuration file from a switch to a TFTP server for storage, follow these steps:

-
- 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 1-10.
 - 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 hostname 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 hostname, 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, see the documentation for your FTP server.

These sections contain this configuration information:

- [Preparing to Download or Upload a Configuration File By Using FTP, page 1-13](#)
- [Downloading a Configuration File By Using FTP, page 1-13](#)
- [Uploading a Configuration File By Using FTP, page 1-15](#)

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, see 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 “Preparing to Download or Upload a Configuration File By Using FTP” section on page 1-13.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	configure terminal	Enter global configuration mode on the switch. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
Step 4	ip ftp username <i>username</i>	(Optional) Change the default remote username.
Step 5	ip ftp password <i>password</i>	(Optional) Change the default password.

	Command	Purpose
Step 6	end	Return to privileged EXEC mode.
Step 7	copy ftp:[[[/[username[:password]@]location]/directory] /filename] system:running-config or copy ftp:[[[/[username[:password]@]location]/directory] /filename] nvram:startup-config	Using FTP, copy the configuration file from a network server to the running configuration or to the startup configuration file.

This example shows how to copy a configuration file named *host1-config* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 and to load and run those commands on the switch:

```
Switch# copy ftp://netadmin1:mypass@172.16.101.101/host1-config system:running-config
Configure using host1-config from 172.16.101.101? [confirm]
Connected to 172.16.101.101
Loading 1112 byte file host1-config:![OK]
Switch#
%SYS-5-CONFIG: Configured from host1-config by ftp from 172.16.101.101
```

This example shows how to specify a remote username of *netadmin1*. The software copies the configuration file *host2-config* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 to the switch startup configuration.

```
Switch# configure terminal
Switch(config)# ip ftp username netadmin1
Switch(config)# ip ftp password mypass
Switch(config)# end
Switch# copy ftp: 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
		Verify that the FTP server is properly configured by referring to the “Preparing to Download or Upload a Configuration File By Using FTP” section on page 1-13.
		Log into the switch through the console port or a Telnet session.
Step 1	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
Step 2	ip ftp username <i>username</i>	(Optional) Change the default remote username.
Step 3	ip ftp password <i>password</i>	(Optional) Change the default password.
Step 4	end	Return to privileged EXEC mode.
Step 5	copy system:running-config ftp:[[//[username[:password]@]location]/directory] /filename] or copy nvram:startup-config ftp:[[//[username[:password]@]location]/directory] /filename]	Using FTP, store the switch running or startup configuration file to the specified location.

This example shows how to copy the running configuration file named *switch2-config* to the *netadmin1* directory on the remote host with an IP address of 172.16.101.101:

```
Switch# copy system:running-config ftp://netadmin1:mypass@172.16.101.101/switch2-config
Write file switch2-config on host 172.16.101.101? [confirm]
Building configuration... [OK]
Connected to 172.16.101.101
Switch#
```

This example shows how to store a startup configuration file on a server by using FTP to copy the file:

```
Switch# configure terminal
Switch(config)# ip ftp username netadmin2
Switch(config)# ip ftp password mypass
Switch(config)# end
Switch# copy nvram:startup-config ftp:
Remote host []? 172.16.101.101
Name of configuration file to write [switch2-config]?
Write file switch2-config on host 172.16.101.101? [confirm]
! [OK]
```

Copying Configuration Files By Using RCP

The 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 hostname.

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.

These sections contain this configuration information:

- [Preparing to Download or Upload a Configuration File By Using RCP, page 1-16](#)
- [Downloading a Configuration File By Using RCP, page 1-17](#)
- [Uploading a Configuration File By Using RCP, page 1-18](#)

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, see 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 “Preparing to Download or Upload a Configuration File By Using RCP” section on page 1-16.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
Step 4	ip rcmd remote-username <i>username</i>	(Optional) Specify the remote username.
Step 5	end	Return to privileged EXEC mode.
Step 6	copy rcp:[[//[username@]location]/directory]/filename] system:running-config or copy rcp:[[//[username@]location]/directory]/filename] nvrn:startup-config	Using RCP, copy the configuration file from a network server to the running configuration or to the startup configuration file.

This example shows how to copy a configuration file named `host1-config` from the `netadmin1` directory on the remote server with an IP address of 172.16.101.101 and load and run those commands on the switch:

```
Switch# copy rcp://netadmin1@172.16.101.101/host1-config system:running-config
Configure using host1-config from 172.16.101.101? [confirm]
Connected to 172.16.101.101
Loading 1112 byte file host1-config:![OK]
Switch#
%SYS-5-CONFIG: Configured from host1-config by rcp from 172.16.101.101
```

This example shows how to specify a remote username of *netadmin1*. Then it copies the configuration file *host2-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: 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 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 “Preparing to Download or Upload a Configuration File By Using RCP” section on page 1-16.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
Step 4	ip rcmd remote-username <i>username</i>	(Optional) Specify the remote username.
Step 5	end	Return to privileged EXEC mode.
Step 6	copy system:running-config rcp:[[//[username@]location]/directory]/filename] or copy nvram:startup-config rcp:[[//[username@]location]/directory]/filename]	Using RCP, copy the configuration file from a switch running or startup configuration file to a network server.

This example shows how to copy the running configuration file named *switch2-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, see the *Cisco IOS Command Reference for Release 12.4*.



Caution

You cannot restore a file after it has been deleted.

Replacing and Rolling Back Configurations

The configuration replacement and rollback feature replaces the running configuration with any saved Cisco IOS configuration file. You can use the rollback function to roll back to a previous configuration.

These sections contain this information:

- [Understanding Configuration Replacement and Rollback, page 1-20](#)
- [Configuration Guidelines, page 1-21](#)
- [Configuring the Configuration Archive, page 1-22](#)
- [Performing a Configuration Replacement or Rollback Operation, page 1-23](#)

Understanding Configuration Replacement and Rollback

- [Archiving a Configuration, page 1-20](#)
- [Replacing a Configuration, page 1-20](#)
- [Rolling Back a Configuration, page 1-21](#)

Archiving a Configuration

The configuration archive provides a mechanism to store, organize, and manage an archive of configuration files. The **configure replace** privileged EXEC command increases the configuration rollback capability. As an alternative, you can save copies of the running configuration by using the **copy running-config destination-url** privileged EXEC command, storing the replacement file either locally or remotely. However, this method lacks any automated file management. The configuration replacement and rollback feature can automatically save copies of the running configuration to the configuration archive.

You use the **archive config** privileged EXEC command to save configurations in the configuration archive by using a standard location and filename prefix that is automatically appended with an incremental version number (and optional timestamp) as each consecutive file is saved. You can specify how many versions of the running configuration are kept in the archive. After the maximum number of files are saved, the oldest file is automatically deleted when the next, most recent file is saved. The **show archive** privileged EXEC command displays information for all the configuration files saved in the configuration archive.

The Cisco IOS configuration archive, in which the configuration files are stored and available for use with the **configure replace** command, is in any of these file systems: FTP, HTTP, RCP, TFTP.

Replacing a Configuration

The **configure replace** privileged EXEC command replaces the running configuration with any saved configuration file. When you enter the **configure replace** command, the running configuration is compared with the specified replacement configuration, and a set of configuration differences is generated. The resulting differences are used to replace the configuration. The configuration replacement operation is usually completed in no more than three passes. To prevent looping behavior no more than five passes are performed.

You can use the **copy source-url running-config** privileged EXEC command to copy a stored configuration file to the running configuration. When using this command as an alternative to the **configure replace target-url** privileged EXEC command, note these major differences:

- The **copy source-url running-config** command is a merge operation and preserves all the commands from both the source file and the running configuration. This command does not remove commands from the running configuration that are not present in the source file. In contrast, the **configure replace target-url** command removes commands from the running configuration that are not present in the replacement file and adds commands to the running configuration that are not present.
- You can use a partial configuration file as the source file for the **copy source-url running-config** command. You must use a complete configuration file as the replacement file for the **configure replace target-url** command.

Rolling Back a Configuration

You can also use the **configure replace** command to roll back changes that were made since the previous configuration was saved. Instead of basing the rollback operation on a specific set of changes that were applied, the configuration rollback capability reverts to a specific configuration based on a saved configuration file.

If you want the configuration rollback capability, you must first save the running configuration before making any configuration changes. Then, after entering configuration changes, you can use that saved configuration file to roll back the changes by using the **configure replace** *target-url* command.

You can specify any saved configuration file as the rollback configuration. You are not limited to a fixed number of rollbacks, as is the case in some rollback models.

Configuration Guidelines

Follow these guidelines when configuring and performing configuration replacement and rollback:

- Make sure that the switch has free memory larger than the combined size of the two configuration files (the running configuration and the saved replacement configuration). Otherwise, the configuration replacement operation fails.
- Make sure that the switch also has sufficient free memory to execute the configuration replacement or rollback configuration commands.
- Certain configuration commands, such as those pertaining to physical components of a networking device (for example, physical interfaces), cannot be added or removed from the running configuration.
 - A configuration replacement operation cannot remove the **interface** *interface-id* command line from the running configuration if that interface is physically present on the device.
 - The **interface** *interface-id* command line cannot be added to the running configuration if no such interface is physically present on the device.
- When using the **configure replace** command, you must specify a saved configuration as the replacement configuration file for the running configuration. The replacement file must be a complete configuration generated by a Cisco IOS device (for example, a configuration generated by the **copy running-config** *destination-url* command).



Note

If you generate the replacement configuration file externally, it must comply with the format of files generated by Cisco IOS devices.

Configuring the Configuration Archive

Using the **configure replace** command with the configuration archive and with the **archive config** command is optional but offers significant benefit for configuration rollback scenarios. Before using the **archive config** command, you must first configure the configuration archive. Starting in privileged EXEC mode, follow these steps to configure the configuration archive:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	archive	Enter archive configuration mode.
Step 3	path <i>url</i>	Specify the location and filename prefix for the files in the configuration archive.
Step 4	maximum <i>number</i>	(Optional) Set the maximum number of archive files of the running configuration to be saved in the configuration archive. <i>number</i> —Maximum files of the running configuration file in the configuration archive. Valid values are from 1 to 14. The default is 10. Note Before using this command, you must first enter the path archive configuration command to specify the location and filename prefix for the files in the configuration archive.
Step 5	time-period <i>minutes</i>	(Optional) Set the time increment for automatically saving an archive file of the running configuration in the configuration archive. <i>minutes</i> —Specify how often, in minutes, to automatically save an archive file of the running configuration in the configuration archive.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify the configuration.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Performing a Configuration Replacement or Rollback Operation

Starting in privileged EXEC mode, follow these steps to replace the running configuration file with a saved configuration file:

	Command	Purpose
Step 1	archive config	(Optional) Save the running configuration file to the configuration archive. Note Enter the path archive configuration command before using this command.
Step 2	configure terminal	Enter global configuration mode.
Step 3		Make necessary changes to the running configuration.
Step 4	exit	Return to privileged EXEC mode.
Step 5	configure replace <i>target-url</i> [list] [force] [time seconds] [no-lock]	Replace the running configuration file with a saved configuration file. <i>target-url</i> —URL (accessible by the file system) of the saved configuration file that is to replace the running configuration, such as the configuration file created in Step 2 by using the archive config privileged EXEC command. list —Display a list of the command entries applied by the software parser during each pass of the configuration replacement operation. The total number of passes also appears. force — Replace the running configuration file with the specified saved configuration file without prompting you for confirmation. time seconds —Specify the time (in seconds) within which you must enter the configure confirm command to confirm replacement of the running configuration file. If you do not enter the configure confirm command within the specified time limit, the configuration replacement operation is automatically stopped. (In other words, the running configuration file is restored to the configuration that existed before you entered the configure replace command). Note You must first enable the configuration archive before you can use the time seconds command line option. no-lock —Disable the locking of the running configuration file that prevents other users from changing the running configuration during a configuration replacement operation.
Step 6	configure confirm	(Optional) Confirm replacement of the running configuration with a saved configuration file. Note Use this command only if the time seconds keyword and argument of the configure replace command are specified.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Working with Software Images

This section describes how to archive (download and upload) software image files, which contain the system software, the Cisco IOS code, and the embedded device manager software.

**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. For switch stacks, the **archive download-sw** and **archive upload-sw** privileged EXEC commands can only be used through the stack master. Software images downloaded to the stack master are automatically downloaded to the rest of the stack members.

To upgrade a switch in the stack that has an incompatible software image, use the **archive copy-sw** privileged EXEC command to copy the software image from an existing stack member to the incompatible switch. That switch automatically reloads and joins the stack as a fully functioning member.

You can download a switch image file from a TFTP, FTP, or RCP server to upgrade the switch software. If you do not have access to a TFTP server, you can download a software image file directly to your PC or workstation by using a web browser (HTTP) and then by using the device manager or Cisco Network Assistant to upgrade your switch. For information about upgrading your switch by using a TFTP server or a web browser (HTTP), see the release notes.

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 to another of the same type.

The protocol that you use depends on which type of server you are using. The FTP and RCP transport mechanisms provide faster performance and more reliable delivery of data than TFTP. These improvements are possible because FTP and RCP are built on and use the TCP/IP stack, which is connection-oriented.

These sections contain this configuration information:

- [Image Location on the Switch, page 1-25](#)
- [tar File Format of Images on a Server or Cisco.com, page 1-25](#)
- [Copying Image Files By Using TFTP, page 1-26](#)
- [Copying Image Files By Using FTP, page 1-29](#)
- [Copying Image Files By Using RCP, page 1-33](#)
- [Copying an Image File from One Stack Member to Another, page 1-38](#)

**Note**

For a list of software images and the supported upgrade paths, see the release notes.

Image Location on the Switch

The Cisco IOS image is stored as a *.bin* file in a directory that shows the version number. A subdirectory contains the files needed for web management. The image is stored on the system board flash memory (flash:).

You can use the **show version** privileged EXEC command to see the software version that is currently running on your switch. In the display, check the line that begins with `System image file is...`. It shows the directory name in flash memory where the image is stored.

You can also use the **dir filesystem:** privileged EXEC command to see the directory names of other software images that might be 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 and web management files

This example shows some of the information contained in the info file. [Table 1-3](#) provides additional details about this information:

```
system_type:0x00000000:image-name
  image_family:xxxx
  stacking_number:x
  info_end:
version_suffix:xxxx
  version_directory:image-name
  image_system_type_id:0x00000000
  image_name:image-nameB.bin
  ios_image_file_size:6398464
  total_image_file_size:8133632
  image_feature:IP|LAYER_3|PLUS|MIN_DRAM_MEG=128
  image_family:xxxx
  stacking_number:x
  board_ids:0x401100c4 0x00000000 0x00000001 0x00000003 0x00000002 0x00008000 0x00008002
0x40110000
  info_end:
```



Note Disregard the `stacking_number` field. It does not apply to the switch.

Table 1-3 *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 and the HTML subdirectory are installed
image_name	Specifies the name of the Cisco IOS image within the tar file
ios_image_file_size	Specifies the Cisco IOS image size in the tar file, which is an approximate measure of how much flash memory is required to hold just the Cisco IOS image

Table 1-3 info File Description (continued)

Field	Description
total_image_file_size	Specifies the size of all the images (the Cisco IOS image and the web management 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_min_dram	Specifies the minimum amount of DRAM needed to run this 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. For switch stacks, the **archive download-sw** and **archive upload-sw** privileged EXEC commands can only be used through the stack master. Software images downloaded to the stack master are automatically downloaded to the rest of the stack members.

To upgrade a switch with an incompatible software image, use the **archive copy-sw** privileged EXEC command to copy the software image from an existing stack member to the incompatible switch. That switch automatically reloads and joins the stack as a fully functioning member.

These sections contain this configuration information:

- [Preparing to Download or Upload an Image File By Using TFTP, page 1-26](#)
- [Downloading an Image File By Using TFTP, page 1-27](#)
- [Uploading an Image File By Using TFTP, page 1-29](#)

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, see 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, go to Step 3.

	Command	Purpose
Step 1		Copy the image to the appropriate TFTP directory on the workstation. Make sure that the TFTP server is properly configured; see the “Preparing to Download or Upload an Image File By Using TFTP” section on page 1-26.
Step 2		Log into the switch through the console port or a Telnet session.

	Command	Purpose
Step 3	archive download-sw /overwrite /reload tftp:[[/location]/directory]/image-name.tar	<p>Download the image file from the TFTP server to the switch, and overwrite the current image.</p> <ul style="list-style-type: none"> • The /overwrite option overwrites the software image in flash memory with the downloaded image. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For <i>/location</i>, specify the IP address of the TFTP server. • For <i>/directory/image-name.tar</i>, specify the directory (optional) and the image to download. Directory and image names are case sensitive.
Step 4	archive download-sw /leave-old-sw /reload tftp:[[/location]/directory]/image-name.tar	<p>Download the image file from the TFTP server to the switch, and keep the current image.</p> <ul style="list-style-type: none"> • The /leave-old-sw option keeps the old software version after a download. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For <i>/location</i>, specify the IP address of the TFTP server. • For <i>/directory/image-name.tar</i>, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it 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 is displayed.

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.

Use the upload feature only if the web management pages associated with the embedded device manager have been installed with the existing image.

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 “Preparing to Download or Upload an Image File By Using TFTP” section on page 1-26.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	archive upload-sw ftp:[[/location]/directory]/image-name.tar	Upload the currently running switch image to the TFTP server. <ul style="list-style-type: none"> For <i>//location</i>, specify the IP address of the TFTP server. For <i>/directory/image-name.tar</i>, specify the directory (optional) and the name of the software image to be uploaded. Directory and image names are case sensitive. The <i>image-name.tar</i> is the name of the software image to be stored on the server.

The **archive upload-sw** privileged EXEC command builds an image file on the server by uploading these files in order: info, the Cisco IOS image, and the web management files. 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. For switch stacks, the **archive download-sw** and **archive upload-sw** privileged EXEC commands can only be used through the stack master. Software images downloaded to the stack master are automatically downloaded to the rest of the stack members.

To upgrade a switch with an incompatible software image, use the **archive copy-sw** privileged EXEC command to copy the software image from an existing stack member to the incompatible switch. That switch automatically reloads and joins the stack as a fully functioning member.

These sections contain this configuration information:

- [Preparing to Download or Upload an Image File By Using FTP, page 1-30](#)
- [Downloading an Image File By Using FTP, page 1-31](#)
- [Uploading an Image File By Using FTP, page 1-32](#)

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 hostname, 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 subnet 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, see 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, go to Step 7.

	Command	Purpose
Step 1		Verify that the FTP server is properly configured by referring to the “Preparing to Download or Upload an Image File By Using FTP” section on page 1-30.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
Step 4	ip ftp username <i>username</i>	(Optional) Change the default remote username.
Step 5	ip ftp password <i>password</i>	(Optional) Change the default password.
Step 6	end	Return to privileged EXEC mode.
Step 7	archive download-sw /overwrite /reload ftp:[[/username[:password]@location]/directory] <i>/image-name.tar</i>	Download the image file from the FTP server to the switch, and overwrite the current image. <ul style="list-style-type: none"> • The /overwrite option overwrites the software image in flash memory with the downloaded image. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For <i>/username[:password]</i>, specify the username and password; these must be associated with an account on the FTP server. For more information, see the “Preparing to Download or Upload an Image File By Using FTP” section on page 1-30. • For <i>@location</i>, specify the IP address of the FTP server. • For <i>directory/image-name.tar</i>, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

Command	Purpose
Step 8 archive download-sw /leave-old-sw /reload ftp:[[/username[:password]@location]/directory] image-name.tar	<p>Download the image file from the FTP server to the switch, and keep the current image.</p> <ul style="list-style-type: none"> • The /leave-old-sw option keeps the old software version after a download. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For //username[:password], specify the username and password. These must be associated with an account on the FTP server. For more information, see the “Preparing to Download or Upload an Image File By Using FTP” section on page 1-30. • For @location, specify the IP address of the FTP server. • For directory/image-name.tar, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it 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 is displayed.

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.

Use the upload feature only if the web management pages associated with the embedded device manager have been installed with the existing image.

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 “Preparing to Download or Upload a Configuration File By Using FTP” section on page 1-13.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
Step 4	ip ftp username <i>username</i>	(Optional) Change the default remote username.
Step 5	ip ftp password <i>password</i>	(Optional) Change the default password.
Step 6	end	Return to privileged EXEC mode.
Step 7	archive upload-sw ftp:[[//[username[:password]@]location]/directory]/image-name.tar	Upload the currently running switch image to the FTP server. <ul style="list-style-type: none"> For <i>//username:password</i>, specify the username and password. These must be associated with an account on the FTP server. For more information, see the “Preparing to Download or Upload an Image File By Using FTP” section on page 1-30. For <i>@location</i>, specify the IP address of the FTP server. For <i>/directory/image-name.tar</i>, specify the directory (optional) and the name of the software image to be uploaded. Directory and image names are case sensitive. The <i>image-name.tar</i> is the name of the software image to be stored on the server.

The **archive upload-sw** command builds an image file on the server by uploading these files in order: info, the Cisco IOS image, and the web management files. 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. For switch stacks, the **archive download-sw** and **archive upload-sw** privileged EXEC commands can only be used through the stack master. Software images downloaded to the stack master are automatically downloaded to the rest of the stack members.

To upgrade a switch with an incompatible software image, use the **archive copy-sw** privileged EXEC command to copy the software image from an existing stack member to the incompatible switch. That switch automatically reloads and joins the stack as a fully functioning member.

These sections contain this configuration information:

- [Preparing to Download or Upload an Image File By Using RCP, page 1-34](#)
- [Downloading an Image File By Using RCP, page 1-35](#)
- [Uploading an Image File By Using RCP, page 1-37](#)

Preparing to Download or Upload an Image File By Using RCP

RCP provides another method of downloading and uploading image files between remote hosts and the switch. Unlike TFTP, which uses User Datagram Protocol (UDP), a connectionless protocol, RCP uses TCP, which is connection-oriented.

To use RCP to copy files, the server from or to which you will be copying files must support RCP. The RCP copy commands rely on the rsh server (or daemon) on the remote system. To copy files by using RCP, you do not need to create a server for file distribution as you do with TFTP. You only need to have access to a server that supports the remote shell (rsh). (Most UNIX systems support rsh.) Because you are copying a file from one place to another, you must have read permission on the source file and write permission on the destination file. If the destination file does not exist, RCP creates it for you.

RCP requires a client to send a remote username on each RCP request to a server. When you copy an image from the switch to a server by using RCP, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **archive download-sw** or **archive upload-sw** privileged EXEC command if a username is specified.
- The username set by the **ip rcmd remote-username *username*** global configuration command if the command is entered.
- The remote username associated with the current TTY (terminal) process. For example, if the user is connected to the router through Telnet and was authenticated through the **username** command, the switch software sends the Telnet username as the remote username.
- The switch hostname.

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, see 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, go to Step 6.

	Command	Purpose
Step 1		Verify that the RCP server is properly configured by referring to the “Preparing to Download or Upload an Image File By Using RCP” section on page 1-34.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
Step 4	ip rcmd remote-username <i>username</i>	(Optional) Specify the remote username.
Step 5	end	Return to privileged EXEC mode.

Command	Purpose
<p>Step 6</p> <p>archive download-sw /overwrite /reload rcp:[[/[username@]location]/directory]/image-name.tar]</p>	<p>Download the image file from the RCP server to the switch, and overwrite the current image.</p> <ul style="list-style-type: none"> • The /overwrite option overwrites the software image in flash memory with the downloaded image. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For <i>//username</i>, specify the username. For the RCP copy request to execute successfully, an account must be defined on the network server for the remote username. For more information, see the “Preparing to Download or Upload an Image File By Using RCP” section on page 1-34. • For <i>@location</i>, specify the IP address of the RCP server. • For <i>/directory/image-name.tar</i>, specify the directory (optional) and the image to download. Directory and image names are case sensitive.
<p>Step 7</p> <p>archive download-sw /leave-old-sw /reload rcp:[[/[username@]location]/directory]/image-name.tar]</p>	<p>Download the image file from the RCP server to the switch, and keep the current image.</p> <ul style="list-style-type: none"> • The /leave-old-sw option keeps the old software version after a download. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For <i>//username</i>, specify the username. For the RCP copy request to execute, an account must be defined on the network server for the remote username. For more information, see the “Preparing to Download or Upload an Image File By Using RCP” section on page 1-34. • For <i>@location</i>, specify the IP address of the RCP server. • For <i>/directory/image-name.tar</i>, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it 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 is displayed.

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.

The upload feature should be used only if the web management pages associated with the embedded device manager have been installed with the existing image.

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 “Preparing to Download or Upload an Image File By Using RCP” section on page 1-34.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
Step 4	ip rcmd remote-username <i>username</i>	(Optional) Specify the remote username.
Step 5	end	Return to privileged EXEC mode.
Step 6	archive upload-sw rcp: [[[// <i>username@</i>] <i>location</i>]/ <i>directory</i>]/ <i>image-name.tar</i>]	Upload the currently running switch image to the RCP server. <ul style="list-style-type: none"> For <i>//username</i>, specify the username; for the RCP copy request to execute, an account must be defined on the network server for the remote username. For more information, see the “Preparing to Download or Upload an Image File By Using RCP” section on page 1-34. For <i>@location</i>, specify the IP address of the RCP server. For <i>/directory</i>]/<i>image-name.tar</i>, specify the directory (optional) and the name of the software image to be uploaded. Directory and image names are case sensitive. The <i>image-name.tar</i> is the name of software image to be stored on the server.

The **archive upload-sw** privileged EXEC command builds an image file on the server by uploading these files in order: info, the Cisco IOS image, and the web management files. 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 an Image File from One Stack Member to Another

For switch stacks, the **archive download-sw** and **archive upload-sw** privileged EXEC commands can be used only through the stack master. Software images downloaded to the stack master are automatically downloaded to the rest of the stack members.

To upgrade a switch that has an incompatible software image, use the **archive copy-sw** privileged EXEC command to copy the software image from an existing stack member to the one that has incompatible software. That switch automatically reloads and joins the stack as a fully functioning member.

**Note**

To successfully use the **archive copy-sw** privileged EXEC command, you must have downloaded from a TFTP server the images for both the stack member switch being added and the stack master. You use the **archive download-sw** privileged EXEC command to perform the download.

Beginning in privileged EXEC mode from the stack member that you want to upgrade, follow these steps to copy the running image file from the flash memory of a different stack member:

	Command	Purpose
Step 1	archive copy-sw / <i>destination-system destination-stack-member-number</i> / force-reload <i>source-stack-member-number</i>	<p>Copy the running image file from a stack member, and then unconditionally reload the updated stack member.</p> <p>Note At least one stack member must be running the image that is to be copied to the switch that is running the incompatible software.</p> <p>For /destination-system <i>destination-stack-member-number</i>, specify the number of the stack member (the destination) to which to copy the source running image file. If you do not specify this stack member number, the default is to copy the running image file to all stack members.</p> <p>Specify /force-reload to unconditionally force a system reload after successfully downloading the software image.</p> <p>For <i>source-stack-member-number</i>, specify the number of the stack member (the source) from which to copy the running image file. The stack member number range is 1 to 9.</p>
Step 2	reload slot <i>stack-member-number</i>	Reset the updated stack member, and put this configuration change into effect.



Unsupported Commands in Cisco IOS Release 12.2(58)SE

This appendix lists some of the command-line interface (CLI) commands that appear when you enter the question mark (?) at the Catalyst 2960 or 2960-S switch prompt but are not supported in this release, either because they are not tested or because of switch hardware limitations. This is not a complete list. These unsupported commands are listed by software feature and command mode:

- [Access Control Lists, page 1-1](#)
- [Boot Loader Commands, page 1-2](#)
- [Debug Commands, page 1-2](#)
- [Embedded Syslog Manager, page 1-2](#)
- [IGMP Snooping Commands, page 1-3](#)
- [Interface Commands, page 1-3](#)
- [MAC Address Commands, page 1-3](#)
- [Miscellaneous, page 1-4](#)
- [Network Address Translation \(NAT\) Commands, page 1-4](#)
- [QoS, page 1-5](#)
- [RADIUS, page 1-5](#)
- [SNMP, page 1-6](#)
- [SNMPv3, page 1-6](#)
- [Spanning Tree, page 1-6](#)
- [VLAN, page 1-6](#)
- [VTP, page 1-7](#)

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*].
show access-lists rate-limit [*destination*]
show accounting
show ip accounting [*checkpoint*] [*output-packets* | *access violations*]
show ip cache [*prefix-mask*] [*type number*]

Unsupported Global Configuration Commands

access-list rate-limit *acl-index* {*precedence* | **mask** *prec-mask*}
access-list dynamic extended

Unsupported Route-Map Configuration Commands

match ip address prefix-list *prefix-list-name* [*prefix-list-name...*]

Boot Loader Commands

Unsupported Global Configuration Commands

boot buffersize

Debug Commands

Unsupported Privileged EXEC Commands

debug platform cli-redirection main
debug platform configuration

Embedded Syslog Manager

Unsupported Global Configuration Commands

All

Unsupported Privileged EXEC Commands

All

IGMP Snooping Commands

Unsupported Global Configuration Commands

`ip igmp snooping tcn`

Interface Commands

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

`transmit-interface type number`

MAC Address Commands

Unsupported Privileged EXEC Commands

`show mac-address-table`

`show mac-address-table address`

`show mac-address-table aging-time`

`show mac-address-table count`

`show mac-address-table dynamic`

`show mac-address-table interface`

`show mac-address-table multicast`

`show mac-address-table notification`

`show mac-address-table static`

`show mac-address-table vlan`

`show mac address-table multicast`



Note Use the **show ip igmp snooping groups** privileged EXEC command to display Layer 2 multicast address-table entries for a VLAN.

Unsupported Global Configuration Commands

mac-address-table aging-time
mac-address-table notification
mac-address-table static

Miscellaneous

Unsupported User EXEC Commands

verify

Unsupported Privileged EXEC Commands

file verify auto
show cable-diagnostics prbs
test cable-diagnostics prbs

Unsupported Global Configuration Commands

errdisable recovery cause unicast flood
l2protocol-tunnel global drop-threshold
memory reserve critical
service compress-config
stack-mac persistent timer

Network Address Translation (NAT) Commands

Unsupported Privileged EXEC Commands

show ip nat statistics
show ip nat translations

QoS

Unsupported Global Configuration Command

priority-list

Unsupported Interface Configuration Commands

ip route-map

priority-group

rate-limit

Unsupported Policy-Map Configuration Command

class class-default where **class-default** is the *class-map-name*.

RADIUS

Unsupported Global Configuration Commands

aaa nas port extended

aaa authentication *feature* default enable

aaa authentication *feature* default line

aaa nas port extended

authentication command bounce-port ignore (only on switches running the LAN Lite image)

authentication command disable-port ignore (only on switches running the LAN Lite image)

radius-server attribute nas-port

radius-server configure

radius-server extended-portnames

SNMP

Unsupported Global Configuration Commands

no monitor session all (only on switches running the LAN Lite image)

snmp-server enable informs

snmp-server enable traps hsrp

snmp-server enable traps rtr (only on switches running the LAN Lite image)

snmp-server ifindex persist

logging discriminator *discr-name* [[**facility**] [**mnemonics**] [**msg-body**] {**drops** *string* | **includes** *string*}] [**severity** {**drops** *sev-num* | **includes** *sev-num*}] [**rate-limit** *msglimit*]

logging buffered discriminator

SNMPv3

Unsupported 3DES Encryption Commands

All

Spanning Tree

Unsupported Global Configuration Command

spanning-tree pathcost method {**long** | **short**}

Unsupported Interface Configuration Command

spanning-tree stack-port

VLAN

Unsupported Global Configuration Command

vlan internal allocation policy {**ascending** | **descending**}

Unsupported vlan-config Command

private-vlan

Unsupported User EXEC Commands

show running-config vlan
show vlan ifindex
vlan database

Unsupported vlan-config Command

private-vlan

Unsupported VLAN Database Commands

vtp
vlan
show vlan private-vlan

VTP

Unsupported Privileged EXEC Commands

vtp {password *password* | pruning | version *number*}



Note

This command has been replaced by the **vtp** global configuration command.



Recommendations for Upgrading a Catalyst 2950 Switch to a Catalyst 2960 Switch

This appendix describes the configuration compatibility issues and the feature behavior differences that you might encounter when you upgrade a Catalyst 2950 switch to a Catalyst 2960 switch.

This appendix consists of these sections:

- [Configuration Compatibility Issues, page 1-1](#)
- [Feature Behavior Incompatibilities, page 1-5](#)

Configuration Compatibility Issues

The configuration commands between the two switch platforms differ for these reasons:

- The Catalyst 2950 switch runs Cisco IOS 12.1EA software, and the Catalyst 2960 switch runs Cisco IOS 12.2SE software.
- The switch families have different hardware.

If you use a Catalyst 2950 switch command, it might not be supported on the Catalyst 2960 switch. The Catalyst 2960 switch software handles the incompatible commands in these ways:

- They are accepted and translated. A message appears.
- They are rejected. A message appears.

In most cases, configuration files are loaded without rejections. [Table 1-1](#) lists the Catalyst 2950 exceptions. The features are listed in alphabetic order, with Catalyst 2950 commands and explanations, and the resulting action on the Catalyst 2960 switch.

Table 1-1 Catalyst 2950 and 2960 Switch Configuration Incompatibilities

Feature	Catalyst 2950 Switch Command and Explanation	Result on the Catalyst 2960 Switch
AAA	<p>These global configuration commands are in Cisco IOS 12.1EA:</p> <pre>aaa preauth aaa processes 1-64 aaa route download 1-1440</pre>	<p>When Cisco IOS 12.2E was restructured, these commands were intentionally removed and are not supported in Cisco IOS 12.2SE.</p> <p>The Catalyst 2960 switch rejects these commands, and this message appears:</p> <pre>Switch(config)# aaa processes 10 ^ %Invalid input detected at '^' marker.</pre>
Clustering	<p>The Catalyst 2950 switch supports only one management VLAN. You can use this global configuration command to change it:</p> <pre>cluster management-vlan vlan-id</pre> <p>This command communicates the management VLAN when the switch is configured for clustering.</p>	<p>With the Catalyst 2960 switch, you can connect to candidate and cluster member switches through any VLAN in common with the cluster command switch.</p> <p>The Catalyst 2960 switch rejects the command, and this message appears:</p> <pre>Switch(config)# cluster management-vlan 2 ^ %Invalid input detected at '^' marker.</pre>
DHCP snooping	<p>A Catalyst 2950 switch DHCP snooping feature limits the number of DHCP packets per second that an interface can receive. You use this interface configuration command to configure it:</p> <pre>ip dhcp snooping limit rate rate</pre> <p>The range is 1 to 4294967294, and by default, the rate limit is not configured.</p>	<p>In Cisco IOS 12.2SE, the range was changed to 1 to 2048 messages per second.</p> <p>The Catalyst 2960 switch accepts any range value. It changes the maximum value to 2048 (if it is more than 2048), and this message appears:</p> <pre>%Invalid input detected at '^' marker.%</pre>
Flow control	<p>The Catalyst 2950 switch supports pause frames on Gigabit Ethernet interfaces. You use this interface configuration command to configure it:</p> <pre>flowcontrol send {desired off on}</pre>	<p>The Catalyst 2960 switch accepts received pause frames but cannot send them. The flowcontrol send command is not supported on the Catalyst 2960 switch.</p> <p>The Catalyst 2960 switch rejects the command, and this message appears:</p> <pre>Switch(config-if)# flowcontrol send desired ^ %Invalid input detected at '^' marker.</pre> <p>You can configure QoS to restrict data traffic without affecting the control traffic. With flow control, all traffic is stopped. For more information, see Chapter 1, "Configuring QoS."</p>

Table 1-1 Catalyst 2950 and 2960 Switch Configuration Incompatibilities (continued)

Feature	Catalyst 2950 Switch Command and Explanation	Result on the Catalyst 2960 Switch
IEEE 802.1x	<p>In Cisco IOS 12.1EA, the Catalyst 2950 switch ranges for the IEEE 802.1x server-timeout, supp-timeout, and tx-period are 1 to 65535. You use these interface configuration commands to configure them:</p> <pre>dot1x timeout server-timeout seconds dot1x timeout supp-timeout seconds dot1x timeout tx-period seconds</pre>	<p>In Cisco IOS 12.2SE, the IEEE 802.1x server-timeout and supp-timeout ranges are 30 to 65535. The tx-period range is 15 to 65535.</p> <p>For server-timeout, the Catalyst 2960 switch accepts 1 to 29 as a valid lower value and changes the value to 30.</p> <p>For supp-timeout, the Catalyst 2960 switch accepts 1 to 29 as a valid lower value and changes the value to 30.</p> <p>For tx-timeout, the Catalyst 2960 switch accepts 1 to 14 as a valid lower value and changes the value to 15.</p> <p>For all three commands, this message appears:</p> <pre>%Invalid input detected at '^' marker.</pre>
IGMP ¹ snooping	<p>The Catalyst 2950 switch implements IGMP snooping based on MAC addresses. You use this global configuration command to configure static groups:</p> <pre>ip igmp snooping vlan vlan-id static mac-address interface interface-id</pre> <p>These Catalyst 2950 switch global configuration commands were implemented to address hardware limitations:</p> <pre>ip igmp snooping source-only-learning [age-timer value] no ip igmp snooping mrouter learn pim v2</pre>	<p>The Catalyst 2960 switch implements IGMP snooping based on IP addresses and uses other advanced hardware. It rejects the Catalyst 2950 IGMP snooping commands, and these messages appear:</p> <pre>Switch(config)# ip igmp snooping vlan 1 static 0002.4b28.c482 interface gigabitethernet0/1 ^ %Invalid input detected at '^' marker. Switch(config)# ip igmp snooping source-only-learning ^ %Invalid input detected at '^' marker. Switch(config)# no ip igmp snooping mrouter learn pim v2 ^ %Invalid input detected at '^' marker.</pre>
Interface MAC address	<p>On the Catalyst 2950 switch, you can set the MAC address for both physical and switch virtual interfaces (SVIs) by using this interface configuration command:</p> <pre>mac-address mac-address</pre>	<p>On the Catalyst 2960 switch, you cannot set the MAC address for physical and SVIs.</p> <p>The switch rejects the command, and this message appears:</p> <pre>Switch(config-if)# mac-address 0100.0ccc.cccc ^ %Invalid input detected at '^' marker.</pre>

Table 1-1 Catalyst 2950 and 2960 Switch Configuration Incompatibilities (continued)

Feature	Catalyst 2950 Switch Command and Explanation	Result on the Catalyst 2960 Switch
QoS ²	<p>There is limited QoS configuration compatibility between the Catalyst 2950 switch and the Catalyst 2960 switch.</p> <p>We recommend that you enable automatic QoS (auto-QoS) on the Catalyst 2950 switch by using the auto qos voip {cisco-phone cisco-softphone trust} interface configuration command.</p> <p>If you have a custom QoS configuration on the Catalyst 2950 switch, we recommend that you use auto-QoS for transition to the Catalyst 2960 switch.</p> <p>Note If auto-QoS does not provide the configuration required for your network, we recommend that you remove the QoS configuration on the Catalyst 2950 switch and create a new configuration on the Catalyst 2960 switch.</p>	<p>The Catalyst 2960 switch accepts the auto qos command and generates QoS commands that are appropriate for the Catalyst 2960 switch. The policer granularity is adjusted to 1 Mbps.</p> <p>For more information about the generated commands, see the auto qos voip command in the command reference for this release.</p>
	<p>Auto-QoS is not enabled on the Catalyst 2950 switch, but other QoS commands are configured.</p>	<p>These Catalyst 2950 switch commands might fail on the Catalyst 2960 switch:</p> <p>mls qos map dscp-cos global configuration command</p> <p>wrr-queue cos-map global configuration command</p> <p>wrr-queue cos-bandwidth global configuration command</p> <p>mls qos trust cos pass-through dscp interface configuration command</p> <p>police policy-map class configuration command</p> <p>The switch might display this message:</p> <pre> ^ %Invalid input detected at '^' marker.</pre>

- RSPAN

The Catalyst 2950 switch uses an extra port, called the reflector port, for its RSPAN implementation. This is not necessary in the Catalyst 2960 switch RSPAN implementation. The Catalyst 2960 switch also supports VLANs as SPAN sources and can forward received packets on SPAN destination ports.

- Multicast

The multicast forwarding decisions on the Catalyst 2960 switch are based on IP addresses. Some Catalyst 2950 switch workarounds to address platform limitations (such as the **ip igmp snooping source-only-learning** global configuration command) are not required on the Catalyst 2960 switch.